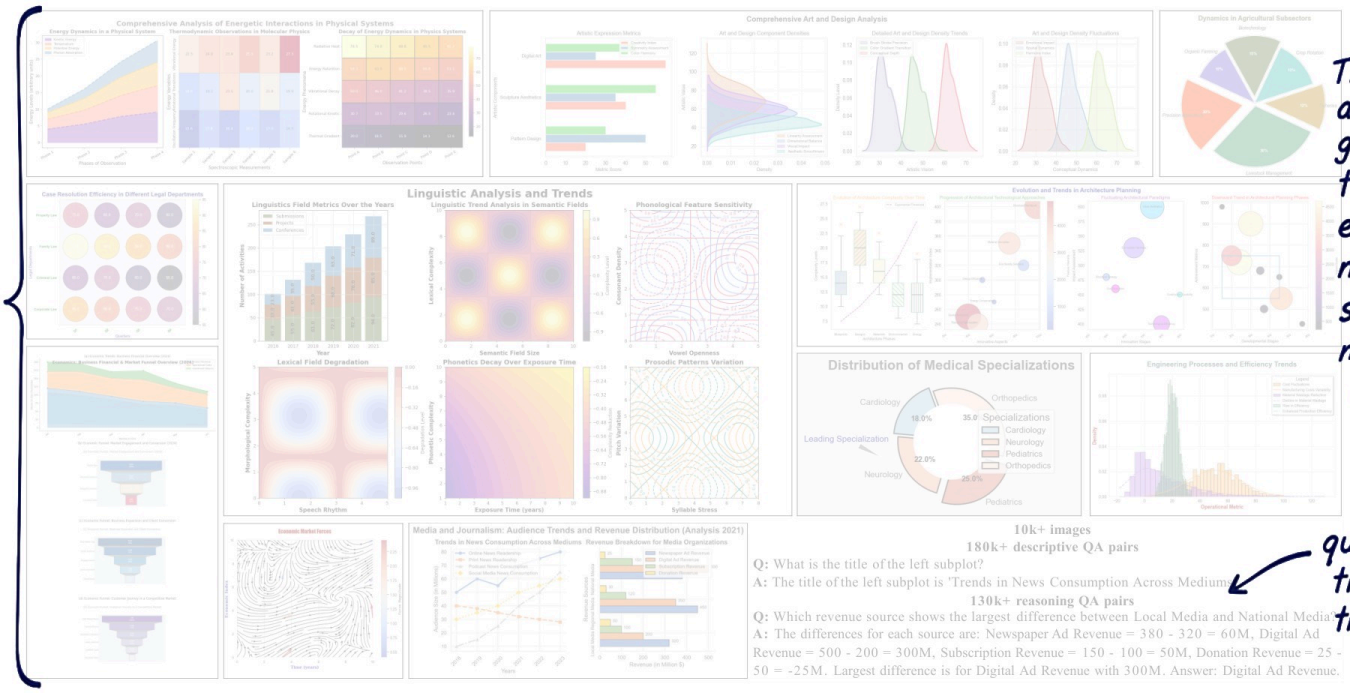
Better synthetic chart training starts with varied plots and grounded questions.

Effective Training Data Synthesis for Improving MLLM Chart Understanding

Yuwei Yang[1], Zeyu Zhang[1], Yunzhong Hou[1], Zhuowan Li[4], Gaowen Liu[3], Ali Payani[3], Yuan-Sen Ting[2], Liang Zheng[1]
[1]Australian National University [2]Ohio State University
[3]Cisco [4]Johns Hopkins University
{yuwei.yang, zeyu.zhang, yunzhong.hou, liang.zheng}@anu.edu.au ting.74@osu.edu
{gaowen.liu, ali.payani}@cisco.com zli110@jhu.edu

arXiv:2508.06492v1 [cs.CV] 8 Aug 2025

29 chart types



The plots are generated training examples, not new scientific measurements.

questions tied to the chart

Figure 1. Overview of the effective chart dataset (ECD). The QA-referenced chart has been adjusted for readability. The ECD is composed of chart images and question-answer (QA) pairs and is intended as a training set for chart understanding. It includes 29 chart types and over 250 subplot combinations, offering more than 10k images and 300k QA pairs with high quality.

Abstract

Being able to effectively read scientific plots, or chart understanding, is a central part toward building effective agents for science. However, existing multimodal large language models (MLLMs), especially open-source ones, are still falling behind with a typical success rate of 30%-50% on challenging benchmarks. Previous studies on fine-tuning MLLMs with synthetic charts are often restricted by their inadequate similarity to the real charts, which could compromise model training and performance on complex real-world charts. In this study, we show that modularizing chart generation and diversifying visual details improves chart understanding capabilities. In particular, we design a five-step data synthesis pipeline, where we separate data and function creation for single plot generation, condition the generation of later subplots on earlier ones for multi-subplot figures, visually diversify the generated figures, filter out low quality data, and finally generate the question-answer (QA) pairs with GPT-4o. This approach allows us to streamline the generation of fine-tuning datasets and introduce the effective chart dataset (ECD), which contains 10k+ chart images and 300k+ QA pairs, covering 25 top-

*This work was supported in part by the Australian Research Council under Discovery Project DP210102801 and Future Fellowship FT240100820, as well as by the National Science Foundation under Grant No. AST-2406729.

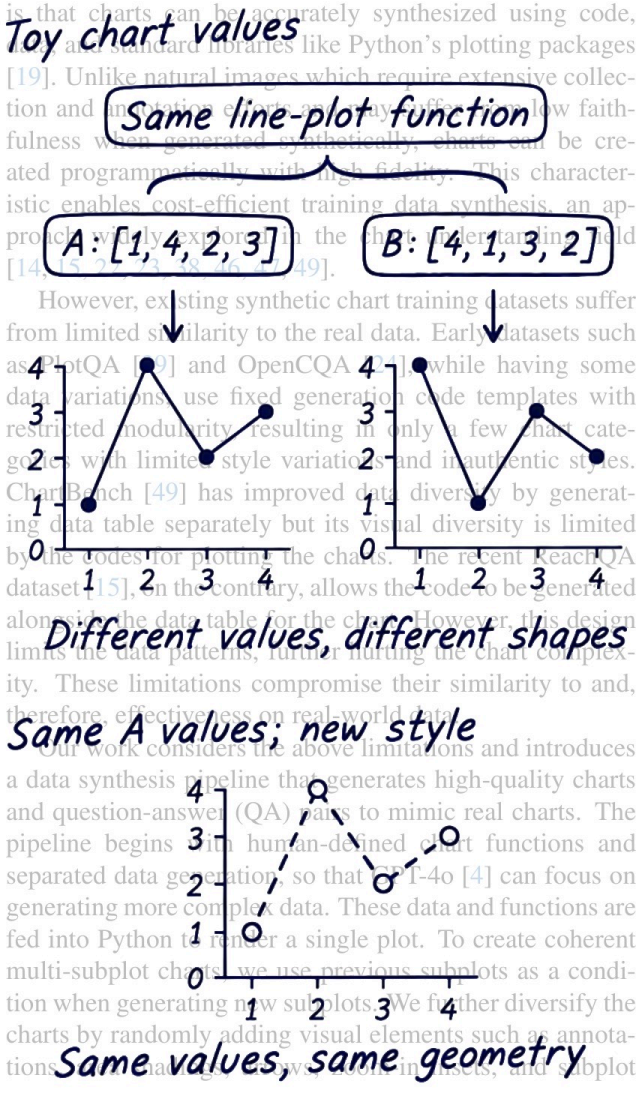
Separate the data from plotting code, then train on what the rendered chart actually shows.

ics and featuring 250+ chart type combinations with high visual complexity. We show that ECD consistently improves the performance of various MLLMs on a range of real-world and synthetic test sets. Code, data and models are available at: https://github.com/yuweiyang-anu/ECD.

1. Introduction

Chart understanding aims to enable MLLMs to accurately answer both descriptive questions (*e.g.*, "what is the label of the x-axis?") that focus on recognizing basic chart elements, and reasoning questions (*e.g.*, "which city demonstrates the second-highest GDP growth rate?") that require analytical capabilities. Several benchmarks extracted from real scientific papers, such as CharXiv [44], have been developed to evaluate model performance on chart understanding tasks. But the performance of open-source models on these benchmarks remains mediocre, highlighting the need for more effective data and training approaches.

An advantage of chart data compared to natural images is that charts can be accurately synthesized using code, using standard libraries like Python's plotting packages [19]. Unlike natural images which require extensive collection and annotation efforts and may suffer from low faithfulness when generated synthetically, charts can be created programmatically with high fidelity. This characteristic enables cost-efficient training data synthesis, an approach widely explored in the chart understanding field [14, 15, 25, 28, 46, 49].



However, existing synthetic chart training datasets suffer from limited similarity to the real data. Early datasets such as PlotQA [39] and OpenCQA [24] while having some data variations, use fixed generation code templates with restricted modularity, resulting in only a few chart categories with limited style variations and inauthentic styles. ChartBench [49] has improved data diversity by generating data table separately but its visual diversity is limited by the codes for plotting the charts. The recent ReachQA dataset [15], on the contrary, allows the code to be generated alongside the data table for the chart. However, this design limits the data patterns, further hurting the chart complexity. These limitations compromise their similarity to and, therefore, effectiveness on real-world data.

Our work considers the above limitations and introduces a data synthesis pipeline that generates high-quality charts and question-answer (QA) pairs to mimic real charts. The pipeline begins with human-defined chart functions and separated data generation, so that GPT-4o [4] can focus on generating more complex data. These data and functions are fed into Python to render a single plot. To create coherent multi-subplot charts, we use previous subplots as a condition when generating new subplots. We further diversify the charts by randomly adding visual elements such as annotations, auxiliary lines, twin axes, insets, and subplot titles. Then, we rate the generated charts based on visual clarity and semantic coherence, filtering out low-quality examples before finally generating QA pairs using GPT-4o.

This pipeline allows us to create the effective chart dataset (ECD). It has over 10,000 images and 300,000+ QA pairs, belonging to 29 chart types and over 250 combinations of different chart types (Fig. 1). By fine-tuning open-source MLLMs with ECD, we show overall improvements on various test sets and that our improvements surpass gains achieved with existing chart training sets.

2. Related Works

The aim is to transfer from synthetic examples to real scientific charts.

Real-world chart datasets such as ChartQA [34], OpenCQA [24], and CharXiv [44] extract charts from online sources, providing genuine examples of visualizations encountered in practice. Despite their authenticity, such real-world datasets suffer from limited scale due to the effort required for collection and annotation.

Synthetic chart datasets have gained popularity due to the programmatic nature of chart generation. Early datasets such as FigureQA [23], DVQA [22], PlotQA [39] and SimChart9k [46] only generate basic chart types (*e.g.*, line, pie, and bar charts) using template-based approaches and exhibit limited complexity. More recent synthetic datasets, including ChartLlama [14], ChartBench [49], ChartAssistant [38], NovaChart [17] and ChartX [47] have expanded the number of images, chart types, and styles, but still demonstrate insufficient visual complexity compared to authentic scientific visualizations. ReachQA [15] uses large language models for code augmentation, but its data complexity and scale are limited. This work presents a new data synthesis pipeline and dataset, and extensive analysis demonstrates that our approach holistically address these limitations.

Multimodal large language models for chart understanding generally adopt two strategies. The first employs a two-stage approach: extracting charts into structured data tables or plotting codes, which then use additional modules for downstream question answering [11, 25, 28, 29, 46]. The second strategy typically leverages chart-caption pairs for visual-language embedding alignment, followed by instruction tuning in an end-to-end manner [10, 14, 18, 27, 30, 35–38, 50, 51, 54, 56, 57, 59]. Given the powerful and generalizable visual-language alignment of MLLMs, performing supervised fine-tuning (SFT) is equally effective.

Data synthesis for training MLLMs is typically classified into two paradigms: (1) collecting real images and then synthesize instructions or (2) jointly synthesizing images and instructions. Some works, such as ProVision [58], LAMM [55], FM^{2}DS [1], DSPT [12] rely on publicly available image datasets or multimodal documents. Structured constructs (*e.g.,* scene graphs) or GPT are then used to generate diverse instructions that emphasize fine-grained visual relationships. To reduce image collection efforts and en-

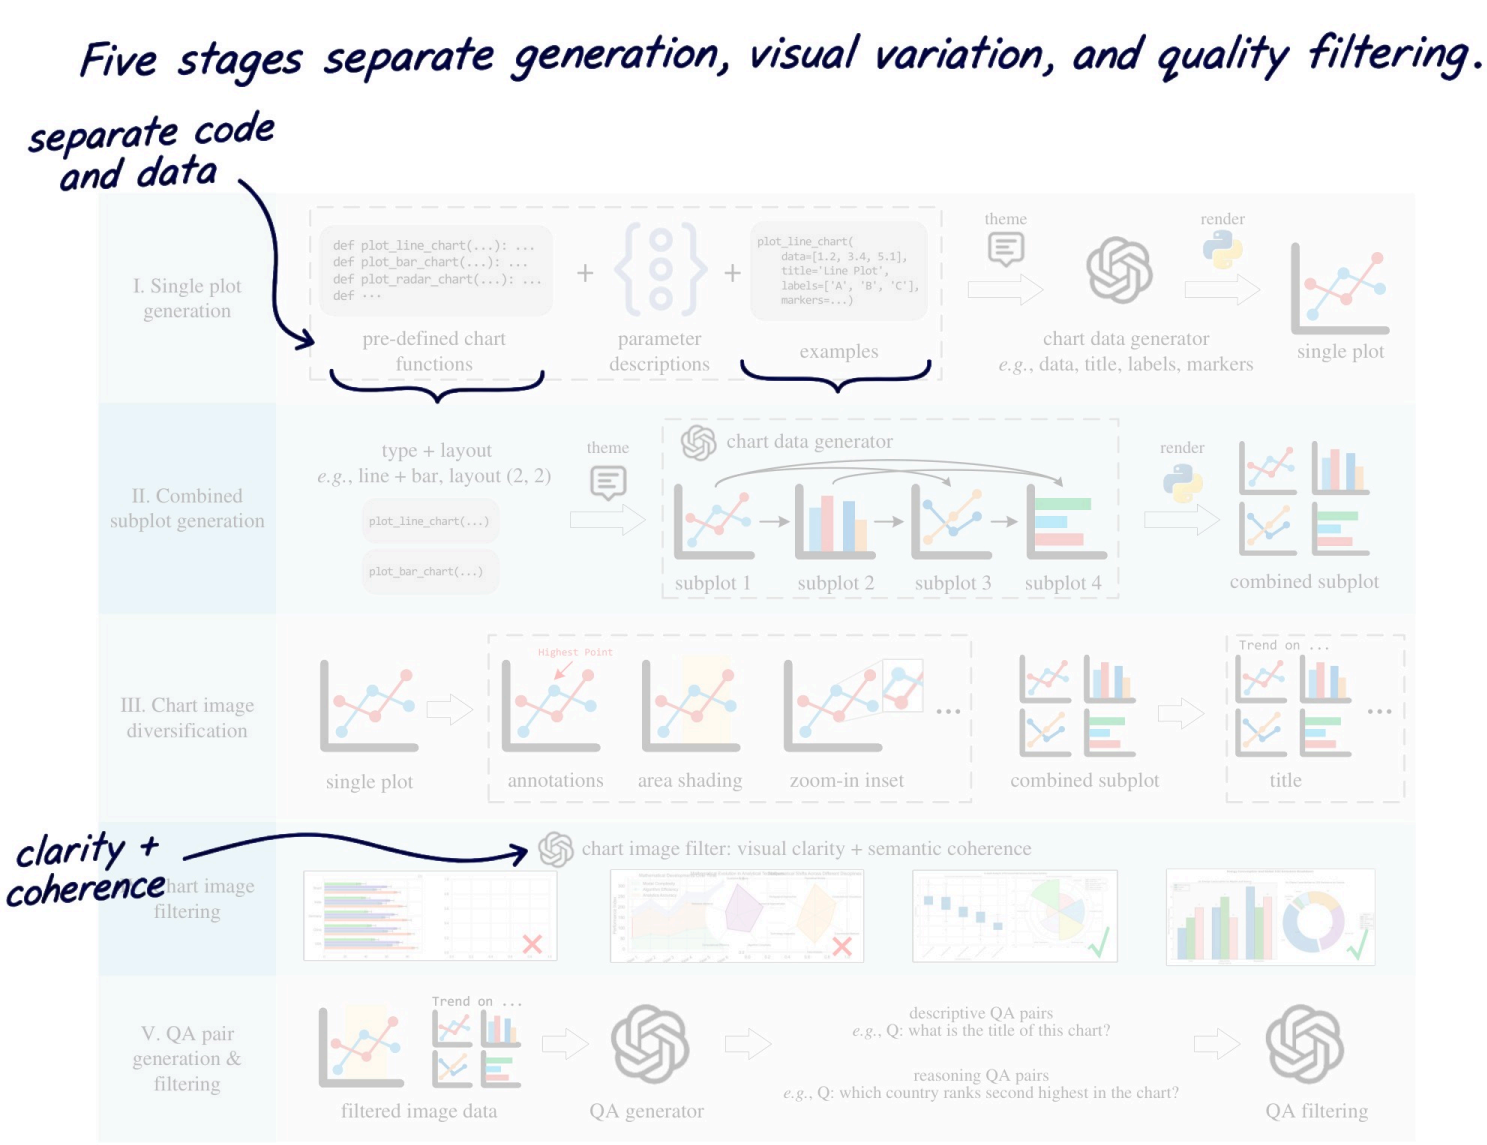


Figure 2. Our chart generation pipeline consists of five key steps: **I. Single plot generation** (Sec. 3.1) creates individual charts using pre-defined chart functions and a separate chart data generator to produce data, titles, labels, and markers for each plot. **II. Combined subplot generation** (Sec. 3.2) produces multi-plot layouts through conditional sequential generation, where each subplot is informed by previous ones to maintain semantic coherence. **III. Chart image diversification** (Sec. 3.3) enhances visual complexity and realism by adding annotations, area shading, zoom-in insets, and other stylistic variations including font modifications and axis border adjustments. **IV. Chart image filtering** (Sec. 3.4) uses rating metrics based on visual clarity and semantic coherence to filter out low-quality chart generations and ensure chart visual quality. **V. QA pair generation & filtering** (Sec. 3.5) creates both descriptive questions and reasoning questions for the chart images, followed by additional quality filtering.

later plots use earlier data

hance privacy protection, a few automatic image-text pair generation methods [3, 9, 21, 32] have been proposed. The aforementioned methods focus on generating natural images. In contrast, programmable images like charts [43, 52], SVGs [41, 48], *etc.* exhibit different characteristics. These images can be fully represented in lossless formats like code, enabling a more precise synthesis of instructions. This paper therefore focuses on chart and QA pair synthesis to enhance MLLMs' chart comprehension capabilities.

3. Proposed Data Synthesis Pipeline

We aim to generate realistic plots whose visual content and stylistic characteristics are similar to real-life scientific plots with the assistance from GPT-4o [4]. However, generating charts with diverse patterns of data while maintaining realism and accuracy of QA pairs presents challenges, and even GPT-4o itself demonstrates imperfections. To address this, we introduce a modular and systematic approach as illustrated in Fig. 2, decomposing the synthesis process into five steps described in the subsections below.

3.1. Generating a Single Plot

The foundation of our pipeline is to generate single plots with rich distributions. Instead of requesting GPT-4o to generate complete visualization codes alongside the data table, we strategically decompose the task by providing GPT with three key inputs: (1) a chart theme (*e.g.*, mathematics, computer science, economics), (2) chart functions that can plot charts of specific types, and (3) parameter descrip-

Filtering removes weak examples after the charts have been diversified.

tions with few-shot examples. Given these inputs, GPT-4o generates appropriate data table and arguments including titles, legends, axis labels, *etc.* By focusing on a step-by-step generation, we ensure the improvement on data distributions and the semantic connection between data values and corresponding textual elements. Specifically, we define 29 chart functions covering distinct chart types including line, bar, pie, scatter, and other common visualization formats. Each function contains modular Python code for rendering a chart based on specific parameters including data, title, axis labels, markers, line styles, *etc.* We generated 10,875 single-plot figures through this process.

3.2. Generating Combined Subplots

To generate combined subplots, we implement a conditional generation approach: given specific chart types and layout configuration, it generates subplots iteratively, conditioned on prior subplots. This ensures consistency across the entire visualization. For example, as illustrated in Fig. 2, when generating subplot 3 (a line chart), our system incorporates data from both subplot 1 (another line chart) and subplot 2 (a bar chart) to maintain thematic coherence. This conditional generation strengthens the relationships between subplots, emulates human-designed scientific figures where multiple subplots present complementary data perspectives. Through this process, we generated 6,006 multi-subplot figures, with an average of 4 subplots per figure.

3.3. Chart Image Diversification

Because of the use of human-defined chart functions in Section 3.1 and Section 3.2, the generated charts exhibit limited style diversity. To solve this problem, we implement a diversification process that modifies the Python code underlying each chart. For single-plot figures, we prompt GPT-4o to alter the code by applying randomly selected visual enhancements from a predefined set of diversification strategies. These include operations such as: *adding arrows*, *annotations*, or *highlights*; *modifying font color*, *style*, or *size*; *using gradient fills* or *area shading*; *removing axis borders*, and *incorporating zoom-in insets* where appropriate.

For multi-subplot figures, we apply diversification after the subplot creation process (Section 3.2), modifying all subplots within one figure through a single GPT-4o request to ensure visual coherence. In addition to the enhancements applied to single-plot figures, we also prompt the model for modifications such as incorporating overarching titles for the entire visualization. Additionally, we enable the use of supplementary visualization libraries beyond Matplotlib, such as Seaborn, to further enhance visual aesthetics. Moreover, we observe that some chart images, especially multi-subplot charts, have overly high resolutions or improper aspect ratios, resulting in unsuitable font sizes. To fix this, we add a post-processing step, prompting GPT-4o to adjust parameters like 'figsize' or 'dpi' in the code for better visualization. We apply diversification techniques and post-processing to all single-plot figures and multi-subplot figures and successfully generate 16,829 figures, substantially increasing the visual complexity and stylistic variation of ECD.

3.4. Filtering Low-quality Charts

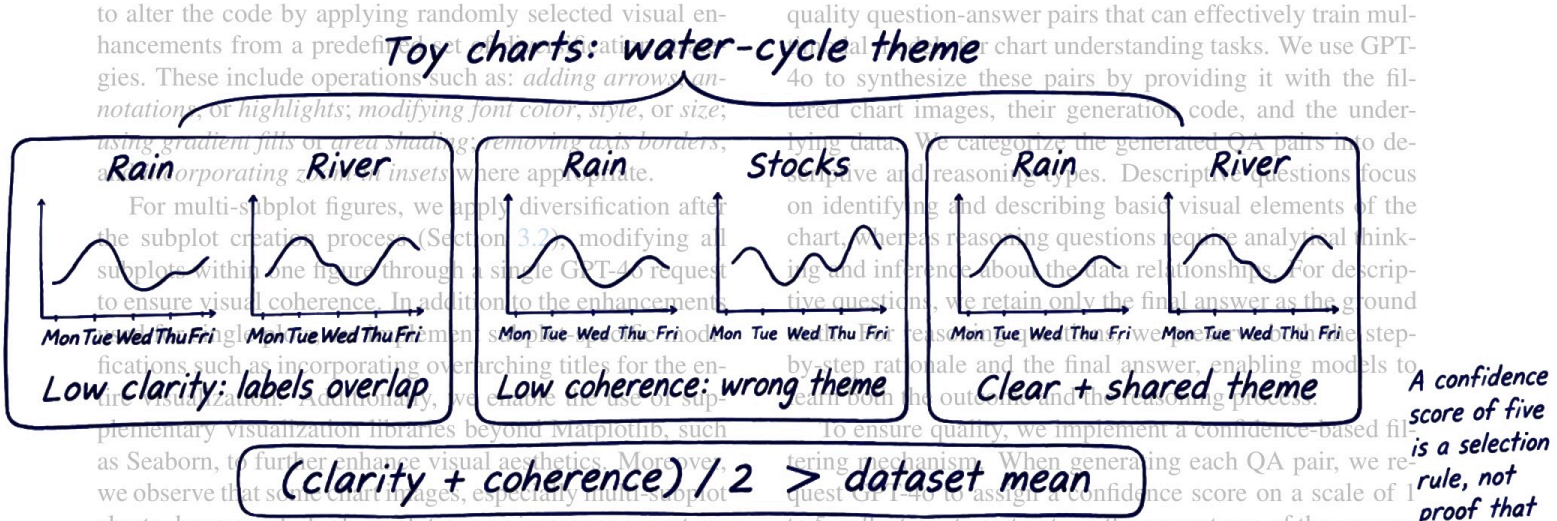
Despite our carefully designed pipeline, we cannot guarantee optimal visual quality for all generated charts. Common issues include excessive blank areas, overcrowded text elements, and misalignment between visual implementation and intended design. To address these challenges, we develop a quality filtering strategy that evaluates charts using two metrics: 1) **visual clarity** $r_{\text{vis}}\left(\boldsymbol{x}, c_{\text{layout}}\right)$, a function of the chart image $\boldsymbol{x}$ and the layout condition c_{layout}; 2) **semantic coherence** $r_{\text{sem}}\left(\boldsymbol{x}, c_{\text{theme}}\right)$, a function of the chart image $\boldsymbol{x}$ and the semantic condition c_{theme}. r_{vis} measures how effectively a chart presents its data, considering readability, appropriate use of visual elements, absence of clutter, *etc.* r_{sem} assesses how well the chart elements and subplots maintain a consistent theme. Both metrics are evaluated using GPT-4o with specialized prompts (see supp. material).

For each chart, we calculate a score by averaging its visual clarity and semantic coherence $r\left(\boldsymbol{x}, c_{\text{layout}}, c_{\text{theme}}\right) = \frac{(r_{\text{s}}+r_{\text{v}})}{2}$, then compare this average to the mean score across the entire dataset $\bar{r}$. We retain only charts whose quality score exceeds the dataset mean, *i.e.*, $r\left(\boldsymbol{x}, c_{\text{layout}}, c_{\text{theme}}\right) > \bar{r}$. This method ensures our final dataset consists of charts that meet or exceed average quality. Starting with 16,829 figures, we filter out approximately 37.4%, resulting in 10,535 high-quality chart images.

3.5. Synthesis and Filtering of QA Pairs

The final stage of our pipeline focuses on generating high-quality question-answer pairs that can effectively train multimodal models for chart understanding tasks. We use GPT-4o to synthesize these pairs by providing it with the filtered chart images, their generation code, and the underlying data. We categorize the generated QA pairs into descriptive and reasoning types. Descriptive questions focus on identifying and describing basic visual elements of the chart, whereas reasoning questions require analytical thinking and inference about the data relationships. For descriptive questions, we retain only the final answer as the ground truth. For reasoning questions, we preserve the whole step-by-step rationale and the final answer, enabling models to learn both the outcome and the reasoning process.

To ensure quality, we implement a confidence-based filtering mechanism. When generating each QA pair, we request GPT-4o to assign a confidence score on a scale of 1 to 5, reflecting its certainty in the correctness of the answer. We only retain pairs with the highest confidence score of 5, eliminating potentially ambiguous or erroneous pairs. Start-



A confidence score of five is a selection rule, not proof that an answer is correct.

16,829 charts → quality filter → 10,535 retained
clarity + coherence

ECD broadens chart variety; visual metrics provide only partial evidence of realism.

Dataset	#Themes	#Chart types	#Type Combo.	#Images	Synthetic	Subplots	#QA Pairs	Training set
OpenCQA [24]	10	5	1	1.2k	×	×	1.2k	×
ChartX [47]	22	18	1	6k	✓	×	6k	×
CharXiv [44]	-	-	-	2.3k	×	✓	5k	×
PlotQA [39]	-	3	1	224k	✓	×	29M	✓
ChartQA [34]	15	3	1	21.9k	×	×	32.7k	✓
ChartLlama [14]	10	10	1	11k	✓	×	160k	✓
ChartBench [49]	9	9	4	68.7k	✓	×	618.5k	✓
ChartAssistant [38]	-	9	1	-	×	×	-	✓
SimChart9k [46]	-	3	1	9k	✓	×	9k	✓
MMC [30]	5	6	-	600k	×	✓	600k	✓
ReachQA [15]	-	10	-	3.7k	✓	✓	22k	✓
ECD (ours)	**25**	**29**	**252**	10.5k	✓	✓	321.5k	✓

Table 1. Comparing the effective chart dataset (ECD) with existing chart datasets. ECD has the most themes, chart types, and type combinations. Although it is not the biggest, experiments show it enables consistent improvements on MLLM chart understanding capabilities.

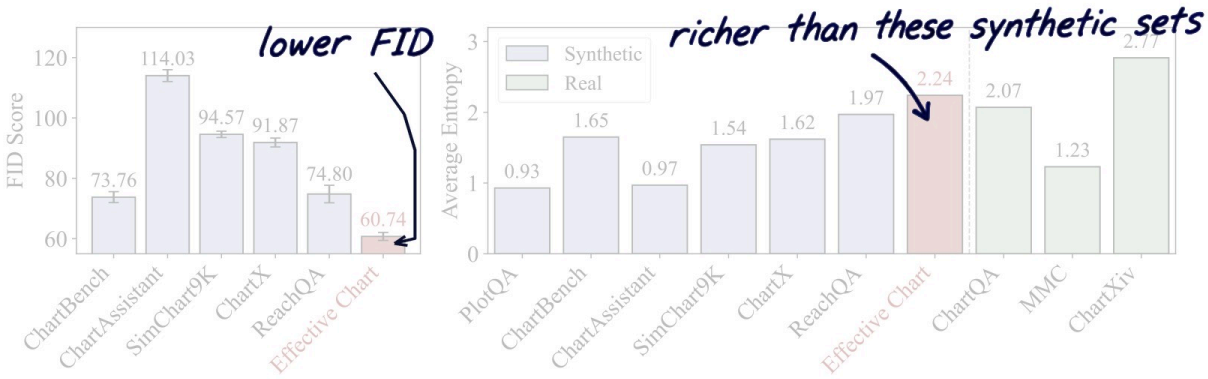


Figure 3. Data complexity and realism comparison beween ECD and other datasets. **Left**: FID comparison. FID is computed between each compared dataset with CharXiv [44]. ECD has the lowest FID score so has stronger alignment with real scientific charts. **Right**: We measure complexity using average entropy of pixel values across images. ECD achieves higher average entropy than the synthetic datasets.

ing with 348,862 initially generated QA pairs from 10,535 images, we filtered out approximately 7.8% based on the confidence scoring mechanism. The final dataset contains 321,544 high-quality QA pairs that cover both descriptive and reasoning tasks.

4. Effective Chart Dataset

The above pipeline allows us to construct the effective chart dataset (ECD). Table 1 presents a comparative analysis between ECD and existing chart datasets.

Styling coverage. ECD has 25 themes and 29 chart types, surpassing all other synthetic training datasets and providing broader contextual coverage across fields including economics, biology, physics, and sociology. It also has 252 combination types, *e.g.*, bar + line and pie + line. This is significantly higher than most datasets that only have single plots, such as PlotQA. CharXiv, MMC, and ReachQA also have some types of combinations, but their counts are unspecified. In terms of scale, ECD contains 10,535 chart images and 321,544 question-answer pairs.

Data realism. We compute Fréchet inception distance (FID) [42] between various datasets and CharXiv, a challenging real-world datasets crawled from arXiv papers. We show the results in Fig. 3 left. We clearly observe that ECD has higher similarity with CharXiv than the other synthetic datasets as measured by FID.

Data complexity. We measure complexity using average entropy (see supp.), computed from image pixel values. For charts, average entropy essentially reflects the sophistication of individual charts, including the number of subplots, line intersections, data points, categorical elements, and the presence of non-trivial data patterns. We show complexity comparison of chart datasets in Fig. 3 right. We observe ECD has high average entropy among synthetic datasets.

ECDBench construction. We construct an additional test set named ECDBench using the same five-stage data generation pipeline described earlier. This process produces single charts (*e.g.*, line charts), multi-layout combination charts with one or two types (*e.g.*, bar charts with a (3, 2) layout, bar + line charts with a (2, 4) layout), and three-type combination charts (*e.g.*, bar + pie + area charts with a (1, 3) layout), resulting in 1,596 chart images after initial filtering.

To ensure data quality, we perform a two-stage rigorous

These metrics describe image appearance, not whether the plotted claims are true.

Four models test transfer across real and synthetic chart benchmarks.

review process. For the first stage, chart image filtering is conducted, where annotators remove images with significant visual recognition issues, yielding 1,224 charts. For minor issues such as overlaps, blank areas, incomplete legends or small fonts, they adjust the code and re-render the charts. In the second stage, annotators carefully verify the visual answerability of the questions by sampling one descriptive and one reasoning-based question from the candidate generated QA pairs for each chart. They also correct inconsistencies within the QA pairs and rewrite overly simplistic, inaccurate, or meaningless QA pairs based on their expertise. The aforementioned process ultimately produces 1,224 high-quality chart images and 2,448 high-quality QA pairs (each accompanied by one descriptive QA and one reasoning-based QA) for subsequent evaluation. More statistical information and evaluations can be found in the supplementary materials.

5. Experiments

5.1. Datasets, MLLMs, and Evaluation Protocols

Datasets. To verify the effectiveness of the effective chart dataset (ECD), we test the trained models across six benchmarks: two real-world datasets (CharXiv [44] and ChartQA [34]) and four synthetic datasets (ReachQA [15], ChartBench [49], ChartX [47] and ECDBench), for a comprehensive evaluation.

Specifically, for real-world data, the recently published **CharXiv** benchmark [44] comprises 2,323 challenging charts extracted from arXiv papers and includes 4,000 descriptive questions about basic chart elements and 1,000 reasoning questions requiring complex visual analysis. **ChartQA** [34] contains charts crawled from 4 online sources, with both human-annotated and machine-generated QA pairs, including a test set of 1,509 images with corresponding 2500 questions. Among the synthetic benchmarks, **ReachQA** [34] includes 500 chart images with 1,000 recognition and 1,000 reasoning questions; **ChartBench** [49] consists of 2,100 images paired with 18,900 QA pairs (2,100 numerical (NQA) and 16,800 binary "yes/no" questions); and **ChartX** [47] contains 1,152 test chart images, each with a corresponding question-answer pair.

MLLMs. We comprehensively evaluate 4 open-source MLLMs: LLaVA-Next-Llama3-8B [26], MiniCPM-V2.6 [53], Phi-3-Vision [2] and the recent Qwen2.5-VL-7B [8], which have diverse architectures and training approaches. We fine-tune these models using ECD with a mix of parameter-efficient and full-parameter fine-tuning methods. We adopt LoRA [16] for LLaVA-Next-Llama3-8B, MiniCPM-V2.6, and Qwen2.5-VL-7B, while performing full-parameter fine-tuning for Phi-3-Vision (4.2B). Specifically, we use a LoRA rank of 8, alpha of 16, a learning rate of 1×10^{-4}, and train for 1 epoch for both LLaVA-Next-Llama3-8B and MiniCPM-V2.6, with batch sizes of 4 (on 4 A100 80G GPUs) and 1 (on a single A100 80G GPU), respectively. Phi-3-Vision is fine-tuned with a batch size of 96 and learning rates of 5×10^{-6} for most parameters and 5×10^{-8} for the projector, while Qwen2.5-VL-7B is fine-tuned with a LoRA rank of 64, alpha of 64, a batch size of 128, and learning rates of 1×10^{-4} for most parameters and 1×10^{-5} for the projector, both trained for 1 epoch on 4 A100 80G GPUs. In all experiments, the vision tower is frozen, and only the remaining parts of models are fine-tuned.

All four vision towers stay frozen; the language-side tuning methods differ.

Evaluation Protocols. We perform QA tests on the official test sets without any specific additional prompts, *e.g.*, chain-of-thought. By default, we use the widely accepted GPT-Acc [15, 44, 47] as the accuracy metric, *i.e.*, using GPT-4o to extract the answer and evaluate the correctness. For ChartQA, ChartBench NQA and ChartX, we follow the official guidelines, allowing a 5% tolerance for questions with numerical answers. Our constructed ECDBench similarly follows the previously established evaluation pattern. For ReachQA and CharXiv, we directly utilize the official prompt templates for testing. Additionally, for the large number of "yes/no" binary questions in ChartBench, we employ a simple regular expression to match the answer and compute the accuracy.

5.2. Main Evaluation

Effective Chart consistently improves MLLMs in chart understanding. We fine-tune the mentioned four open-source MLLMs with ECD and evaluate the fine-tuned models on six test sets. Results are summarized in Tab. 2. We observe that ECD improves the overall performance of four MLLMs. For instance, LLaVA-Next-Llama3-8B's accuracy increases from 64.56% to 68.64% and from 27.69% to 46.61% on ChartQA and ChartX, respectively. On the challenging CharXiv benchmark, ECD boosts Phi-3-Vision's accuracy from 31.50% and 60.52% to 33.40% and 68.00% for reasoning and descriptive questions.

ECD fine-tuning causes slight performance drops in specific metrics, as shown by the red-highlighted results in Tab. 2. MiniCPM-V2.6 and Qwen2.5-VL-7B see decreases on "Binary" questions in ChartBench, likely due to its unique distribution with "yes/no" questions only. For the latest SOTA Qwen2.5-VL-7B, while it achieves notable average improvements on CharXiv and ReachQA, minor drops (*e.g.*, reasoning accuracy on CharXiv, from 41.20% to 40.20%, recognition accuracy on ReachQA, from 71.90% to 70.80%) may be attributed to its pre-training on clean chart data, which closely resembles certain benchmarks. Moreover, the data mixture ratio during fine-tuning might not have been fully optimized for specific tasks or models.

On ECDBench, performance significantly improves due to the similar QA distribution advantage: LLaVA-Next-

ECD training → 4 MLLMs → 2 real-chart benchmarks; 4 synthetic benchmarks

Overall gains coexist with specific regressions; read the individual tasks as well as averages.

Models	CharXiv			ChartQA	ReachQA			ChartBench			ChartX	ECDBench		
	Rea.	Des.	Avg.	QA	Rea.	Recog.	Avg.	Binary	NQA	Avg.	QA	Rea.	Des.	Avg.
Baselines														
Human	80.50	92.10	89.78	-	65.10	84.60	74.85	-	-	-	-	-	-	-
Random (GPT-4o) [4]	10.80	19.85	18.04	30.04	8.20	13.30	10.75	-	22.73	-	11.46	4.58	1.63	3.10
Proprietary multimodal large language models														
GPT-4o mini [4]	34.10	74.92	66.76	77.52	27.20	53.50	40.35	-	34.93	-	44.36	24.26	57.27	40.77
GPT-4o [4]	47.10	84.45	76.98	85.70	39.70	66.80	53.25	-	52.88	-	58.33	35.62	70.18	52.90
Claude 3.5 Sonnet [5]	60.20	84.30	79.48	90.80	51.70	74.30	63.00	-	48.29	-	42.71	41.99	68.14	55.07
Open-source multimodal large language models														
LLaVA-Next-Llama3-8B [26]	19.70	38.90	35.06	64.56	8.00	23.30	15.65	61.18	38.29	58.64	27.69	4.74	17.16	10.95
+ ECD	21.80	59.05	51.60	68.64	12.70	37.50	25.10	64.21	47.05	62.30	46.61	16.50	46.65	31.58
MiniCPM-V2.6 [53]	25.40	52.65	47.20	76.96	14.50	48.10	31.30	75.67	52.33	73.08	47.31	15.15	39.95	27.53
+ ECD	26.90	61.90	54.91	77.28	17.50	48.40	32.95	70.77	52.81	68.77	50.43	18.14	52.21	35.17
Phi-3-Vision [2]	31.50	60.52	54.72	81.92	28.40	62.80	45.60	71.92	62.71	70.90	67.53	21.65	41.18	31.41
+ ECD	33.40	68.00	61.08	84.88	32.80	64.50	48.65	73.49	63.76	72.41	71.44	29.49	59.31	44.40
Qwen2.5-VL-7B [8]	41.20	66.40	61.36	83.04	29.70	71.90	50.80	80.99	67.81	79.53	67.80	19.04	57.35	38.19
+ ECD	40.20	74.20	67.40	85.32	36.20	70.80	53.50	79.35	70.86	78.41	70.83	35.38	66.34	50.86

35.06 to 51.60

41.20 to 40.20

Table 2. Comparing MLLMs on six test sets. "Rea.", "Des.", and "Recog." mean reasoning, descriptive, and recognition questions, respectively. "Binary" denotes "yes/no" answers, and NQA means numerical QAs. When fine-tuned on our effective chart dataset, except a few individual metrics (highlighted in red), the four open-source MLLMs display overall improvements (highlighted in green).

In this LLaVA comparison, ECD leads four of the six test-set columns.

Training set	Test set					
	CharXiv	ChartQA	ReachQA	ChartBench	ChartX	ECDBench
Baseline	35.06	64.56	15.65	58.64	27.69	10.95
ChartQA	35.16	**68.92**	15.00	61.41	31.51	13.11
ChartBench	32.86	61.56	18.35	**75.13**	37.33	10.99
ReachQA	30.68	64.50	24.35	57.59	39.24	13.48
ECD	**51.60**	68.64	**25.10**	62.30	**46.61**	**31.58**

Table 3. Comparing different chart training sets of their fine-tuning performance on 6 test sets. Green indicates performance growth and red highlights decline. We use LLaVA-Next-Llama3-8B [26] as the baseline and fine-tune it on ChartQA [34], ChartBench [49], ReachQA [15], and ECD. For each column, best number is in **bold**, and second best underlined.

Llama3-8B's average accuracy increases from 10.95% to 31.58%, and Qwen2.5-VL-7B achieves the highest accuracy, rising from 38.19% to 50.86%.

Comparison with other training datasets. We then compare ECD with other chart understanding training sets, including ChartQA [34], ChartBench [49], and ReachQA [15]. To do so, we use these datasets to fine-tune LLaVA-Next-Llama3-8B, evaluate the fine-tuned models on six test sets, and report the average accuracy. Results are presented in Tab. 3.

While the existing training sets improve on their own test sets and those with a similar distribution, for other test sets, the fine-tuned model often results in performance drops. For example, since ChartBench has a very different QA distribution, comprising mostly binary QA, training on it gives the best evaluation result for this dataset. However, it also falls behind on two other benchmarks.

In comparison, ECD consistently improves the baseline performance across 6 test sets, while other chart training datasets show the mixed and less consistent results. In addition, ECD yields the largest improvement on the CharXiv, ReachQA, ChartX, and our ECDBench, and the second-largest improvement on ChartQA and ChartBench.

5.3. Ablation Study

To isolate the effects of specific design choices, we conduct a series of ablation studies. When comparing variants of ECD in this section, we maintain consistent image counts for fair comparisons. We select LLaVA-Next-Llama3-8B [26] as the base model for our study and conduct experiments on each plot image by randomly selecting five QA pairs (including both descriptive and reasoning QA pairs) to examine the chart types, themes, and QA type.

Impact of number of chart types. As shown in Fig. 4 (a), increasing the number of chart types from 5 to 29 yields progressive improvements on the CharXiv test set. This confirms the importance of diverse training chart types for effective generalization to real-world charts.

Impact of number of chart themes. Fig. 4 (b) demonstrates that expanding thematic coverage improves model performance, though the improvement gradually slows as the number of themes increases, with the best results achieved at 25 themes. This indicates the usefulness of having more themes.

Impact of QA Type. Fig. 4 (c) compares the pre-trained LLaVA-Next baseline and with the same fine-tuned with three ECD variants, including one with descriptive questions only, one with reasoning questions only, and the full ECD. We observe that using either descriptive (+4.16%) or reasoning questions (+6.40%) improves performance, with greater gains observed when using reasoning questions exclusively. The best model is obtained with the full ECD fine-tuning, which yields 47.02% accuracy. Results validate the inclusion of both types of questions.

Mixed questions help, while adding more examples can reach a plateau.

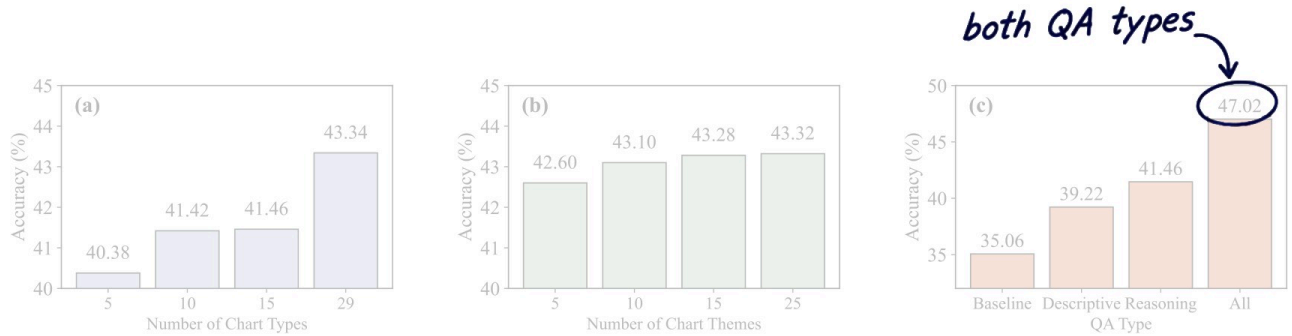


Figure 4. Impact of the number of chart types (a), chart themes (b), and QA type (c). Average accuracy on CharXiv is reported. In the compared variants, we make sure the numbers of images are the same. LLaVA-Next-Llama3-8B is selected as the baseline model.

5.4. Further Analysis

Impact of data scale. Experimental results on data scaling are shown in Tab. 4. We find that increasing data size (default scale is 10k) consistently improves average performance on the synthetic benchmark ReachQA. For real-world benchmark CharXiv, the gains saturate beyond 20k, probably due to the complexity of this task. As future work, we will scale up data in the pre-training phase.

Test set	2k	5k	10k	20k	30k	40k
CharXiv	40.46	45.66	46.84	48.90	47.84	48.26
ReachQA	18.25	18.85	20.85	22.25	23.50	24.75

not monotonic

Table 4. Impact of data scale (5 selected QA pairs per image).

Impact of data mixture ratios. We tried different ratios of descriptive and reasoning questions of ECD (50k QA pairs in total). Results on our CharXiv are shown in Tab. 5. Interestingly, we find the average accuracy metrics peak at a 2:3 or 1:1 ratio. This may be because our definition of descriptive/reasoning does not perfectly match the definition on the benchmarks. Another possible reason is the benefit of having both types of questions on VLM training, which needs further exploration.

CharXiv	0:5	1:4	2:3	1:1	3:2	4:1	5:0
Descriptive Acc	51.25	53.88	53.75	54.90	54.25	52.92	47.48
Reasoning Acc	20.20	19.20	20.90	20.50	20.00	20.90	22.50
Average Acc	45.04	46.94	47.18	48.02	47.40	46.52	42.48

Table 5. Impact of data ratios (descriptive vs. reasoning).

Impact of visual diversification. We use FID (with CharXiv) and average pixel entropy to evaluate image quality with or without the diversification step (Section 3.3). Results are presented in Tab. 6 (a) and (b). It is shown that adding diversification is very beneficial, which decreases FID by 19.64 and increases average entropy by 0.57.

Importance of image and QA quality filtering. In Tab. 6 (c) and (d), we compare dataset quality with or without image quality filtering (Section 3.4). If we remove quality filtering, the FID score with the CharXiv dataset increases by 5.46, and the average entropy decreases by 0.08. The results show that filtering out low-quality images is beneficial. Moreover, to validate the effectiveness of QA filtering (Section 3.5), we remove the QA filtering and compare the fine-tuned LLaVA-Next-Llama3-8B. Results are shown in Tab. 7. We observe that ECD without QA filtering still improves the baseline by 11.58% and that with QA filtering the improvement becomes 11.96%. The additional improvement outlines the benefit of doing so.

Data Strategy	FID↓	Avg. Entropy↑
(a) w/o Visual Diversification	80.38 ± 1.45	1.67
(b) w Visual Diversification	**60.74 ± 1.32**	**2.24**
(c) w/o Image Filtering	66.20 ± 1.11	2.16
(d) w Image Filtering	**60.74 ± 1.32**	**2.24**

Table 6. Comparing data processing strategies. We compare visual diversification and the image filtering strategies with not applying them. We report FID (with CharXiv) and average pixel entropy.

	CharXiv		
	Rea.	Des.	Avg.
LLaVA-Next-Llama3-8B	19.70	38.90	35.06
w/o QA Filtering	19.50	53.42	46.64
w QA Filtering	**20.90**	**53.55**	**47.02**

QA filtering adds 0.38 percentage points here: 46.64 to 47.02.

Table 7. Quantitative performance comparison before and after QA filtering. The average performance improves after filtering.

6. Conclusion

This paper introduces a chart synthesis pipeline and the effective chart dataset (ECD), a high-quality multimodal training resource that enhances chart understanding capabilities in existing MLLMs. Through systematic ablation studies and comparative evaluations, we demonstrate how each component of our pipeline contributes to the overall effectiveness of the resulting dataset. The effective chart dataset generally improves the accuracy of MLLMs across multiple test sets. Particularly notable is the improvement on challenging real-world chart benchmarks like CharXiv, highlighting our approach's ability to bridge the gap between synthetic training data and authentic scientific visualizations. The modular nature of our pipeline enables further extensions and customizations, while our findings regarding the importance of quality and complexity in synthetic data generation have broader implications for multimodal training dataset development.

References

[1] Amirhossein Abaskohi, Spandana Gella, Giuseppe Carenini, and Issam H Laradji. Fm2ds: Few-shot multimodal multihop data synthesis with knowledge distillation for question answering. *arXiv preprint arXiv:2412.07030*, 2024. 2

[2] Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, et al. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*, 2024. 6, 7

[3] Hossein Aboutalebi, Hwanjun Song, Yusheng Xie, Arshit Gupta, Justin Sun, Hang Su, Igor Shalyminov, Nikolaos Pappas, Siffi Singh, and Saab Mansour. Magid: An automated pipeline for generating synthetic multi-modal datasets. *arXiv preprint arXiv:2403.03194*, 2024. 3

[4] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 2, 3, 7, 12

[5] Anthropic. Introducing claude 3.5 sonnet. https://www.anthropic.com/news/claude-3-5-sonnet, 2024. 7, 12

[6] Anthropic. Claude 3.7 sonnet and claude code. https://www.anthropic.com/news/claude-3-7-sonnet, 2025. 12

[7] Anthropic. Introducing claude 4. https://www.anthropic.com/news/claude-4, 2025. 12

[8] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025. 6, 7, 12

[9] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18392–18402, 2023. 3

[10] Victor Cărbune, Hassan Mansoor, Fangyu Liu, Rahul Aralikatte, Gilles Baechler, Jindong Chen, and Abhanshu Sharma. Chart-based reasoning: Transferring capabilities from llms to vlms. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 989–1004, 2024. 2

[11] Jinyue Chen, Lingyu Kong, Haoran Wei, Chenglong Liu, Zheng Ge, Liang Zhao, Jianjian Sun, Chunrui Han, and Xiangyu Zhang. Onechart: Purify the chart structural extraction via one auxiliary token. In *Proceedings of the 32nd ACM International Conference on Multimedia (ACM MM)*, pages 147–155, 2024. 2

[12] Daixuan Cheng, Shaohan Huang, Ziyu Zhu, Xintong Zhang, Wayne Xin Zhao, Zhongzhi Luan, Bo Dai, and Zhenliang Zhang. On domain-specific post-training for multimodal large language models. *arXiv preprint arXiv:2411.19930*, 2024. 2

[13] Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025. 12

[14] 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, 5

[15] Wei He, Zhiheng Xi, Wanxu Zhao, Xiaoran Fan, Yiwen Ding, Zifei Shan, Tao Gui, Qi Zhang, and Xuanjing Huang. Distill visual chart reasoning ability from llms to mllms. *arXiv preprint arXiv:2410.18798*, 2024. 2, 5, 6, 7, 14, 17

[16] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022. 6

[17] Linmei Hu, Duokang Wang, Yiming Pan, Jifan Yu, Yingxia Shao, Chong Feng, and Liqiang Nie. Novachart: A large-scale dataset towards chart understanding and generation of multimodal large language models. In *Proceedings of the 32nd ACM International Conference on Multimedia (ACM MM)*, pages 3917–3925, 2024. 2

[18] Muye Huang, Lai Han, Xinyu Zhang, Wenjun Wu, Jie Ma, Lingling Zhang, and Jun Liu. Evochart: A benchmark and a self-training approach towards real-world chart understanding. *arXiv preprint arXiv:2409.01577*, 2024. 2

[19] J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007. 2

[20] Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card. *arXiv preprint arXiv:2412.16720*, 2024. 12

[21] Qirui Jiao, Daoyuan Chen, Yilun Huang, Yaliang Li, and Ying Shen. Img-diff: Contrastive data synthesis for multimodal large language models. *arXiv preprint arXiv:2408.04594*, 2024. 3

[22] Kushal Kafle, Brian Price, Scott Cohen, and Christopher Kanan. Dvqa: Understanding data visualizations via question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5648–5656, 2018. 2

[23] Samira Ebrahimi Kahou, Vincent Michalski, Adam Atkinson, Ákos Kádár, Adam Trischler, and Yoshua Bengio. Figureqa: An annotated figure dataset for visual reasoning. *arXiv preprint arXiv:1710.07300*, 2017. 2

[24] Shankar Kantharaj, Xuan Long Do, Rixie Tiffany Leong, Jia Qing Tan, Enamul Hoque, and Shafiq Joty. Opencqa: Open-ended question answering with charts. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 11817–11837, 2022. 2, 5

[25] Kenton Lee, Mandar Joshi, Iulia Raluca Turc, Hexiang Hu, Fangyu Liu, Julian Martin Eisenschlos, Urvashi Khandelwal, Peter Shaw, Ming-Wei Chang, and Kristina Toutanova. Pix2struct: Screenshot parsing as pretraining for visual language understanding. In *International Conference on Machine Learning (ICML)*, pages 18893–18912. PMLR, 2023. 2

[26] Bo Li, Kaichen Zhang, Hao Zhang, Dong Guo, Renrui Zhang, Feng Li, Yuanhan Zhang, Ziwei Liu, and Chunyuan

Li. Llava-next: Stronger llms supercharge multimodal capabilities in the wild (2024). *URL https://llava-vl. github. io/blog/2024-05-10-llava-next-stronger-llms*. 6, 7

[27] Zhuowan Li, Bhavan Jasani, Peng Tang, and Shabnam Ghadar. Synthesize step-by-step: Tools templates and llms as data generators for reasoning-based chart vqa. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13613–13623, 2024. 2

[28] Fangyu Liu, Julian Eisenschlos, Francesco Piccinno, Syrine Krichene, Chenxi Pang, Kenton Lee, Mandar Joshi, Wenhu Chen, Nigel Collier, and Yasemin Altun. Deplot: One-shot visual language reasoning by plot-to-table translation. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 10381–10399, 2023. 2

[29] Fangyu Liu, Francesco Piccinno, Syrine Krichene, Chenxi Pang, Kenton Lee, Mandar Joshi, Yasemin Altun, Nigel Collier, and Julian Eisenschlos. Matcha: Enhancing visual language pretraining with math reasoning and chart derendering. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12756–12770, 2023. 2

[30] Fuxiao Liu, Xiaoyang Wang, Wenlin Yao, Jianshu Chen, Kaiqiang Song, Sangwoo Cho, Yaser Yacoob, and Dong Yu. Mmc: Advancing multimodal chart understanding with large-scale instruction tuning. *arXiv preprint arXiv:2311.10774*, 2023. 2, 5

[31] Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. Mmbench: Is your multi-modal model an all-around player? In *European conference on computer vision (ECCV)*, pages 216–233. Springer, 2024. 13

[32] Zheng Liu, Hao Liang, Xijie Huang, Wentao Xiong, Qinhan Yu, Linzhuang Sun, Chong Chen, Conghui He, Bin Cui, and Wentao Zhang. Synthvlm: High-efficiency and high-quality synthetic data for vision language models. *arXiv preprint arXiv:2407.20756*, 2024. 3

[33] Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts. *arXiv preprint arXiv:2310.02255*, 2023. 13

[34] Ahmed Masry, Xuan Long Do, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. Chartqa: A benchmark for question answering about charts with visual and logical reasoning. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2263–2279, 2022. 2, 5, 6, 7, 14, 17

[35] Ahmed Masry, Parsa Kavehzadeh, Xuan Long Do, Enamul Hoque, and Shafiq Joty. Unichart: A universal vision-language pretrained model for chart comprehension and reasoning. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 14662–14684, 2023. 2

[36] Ahmed Masry, Mehrad Shahmohammadi, Md Rizwan Parvez, Enamul Hoque, and Shafiq Joty. Chartinstruct: Instruction tuning for chart comprehension and reasoning. *arXiv preprint arXiv:2403.09028*, 2024.

[37] Ahmed Masry, Megh Thakkar, Aayush Bajaj, Aaryaman Kartha, Enamul Hoque, and Shafiq Joty. Chartgemma: Visual instruction-tuning for chart reasoning in the wild. *arXiv preprint arXiv:2407.04172*, 2024.

[38] Fanqing Meng, Wenqi Shao, Quanfeng Lu, Peng Gao, Kaipeng Zhang, Yu Qiao, and Ping Luo. Chartassisstant: A universal chart multimodal language model via chart-to-table pre-training and multitask instruction tuning. *arXiv preprint arXiv:2401.02384*, 2024. 2, 5, 14, 17

[39] Nitesh Methani, Pritha Ganguly, Mitesh M Khapra, and Pratyush Kumar. Plotqa: Reasoning over scientific plots. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 1527–1536, 2020. 2, 5, 14, 17

[40] OpenAI. Openai o3 and o4-mini system card. https://cdn.openai.com/pdf/2221c875-02dc-4789-800b-e7758f3722c1/o3-and-o4-mini-system-card.pdf, 2025. 12

[41] Juan A Rodriguez, Shubham Agarwal, Issam H Laradji, Pau Rodriguez, David Vazquez, Christopher Pal, and Marco Pedersoli. Starvector: Generating scalable vector graphics code from images. *arXiv preprint arXiv:2312.11556*, 2023. 3

[42] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. *Advances in neural information processing systems*, 29, 2016. 5, 14

[43] Chufan Shi, Cheng Yang, 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. *arXiv preprint arXiv:2406.09961*, 2024. 3

[44] Zirui Wang, Mengzhou Xia, Luxi He, Howard Chen, Yitao Liu, Richard Zhu, Kaiqu Liang, Xindi Wu, Haotian Liu, Sadhika Malladi, et al. Charxiv: Charting gaps in realistic chart understanding in multimodal llms. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 2, 5, 6

[45] xAI Organization. Realworldqa: A benchmark for real-world understanding. https://huggingface.co/datasets/xai-org/RealworldQA, 2025. 13

[46] Renqiu Xia, Bo Zhang, Haoyang Peng, Hancheng Ye, Xiangchao Yan, Peng Ye, Botian Shi, Yu Qiao, and Junchi Yan. Structchart: Perception, structuring, reasoning for visual chart understanding. *arXiv preprint arXiv:2309.11268*, 2023. 2, 5, 14, 17

[47] Renqiu Xia, Bo Zhang, Hancheng Ye, Xiangchao Yan, Qi Liu, Hongbin Zhou, Zijun Chen, Min Dou, Botian Shi, Junchi Yan, et al. Chartx & chartvlm: A versatile benchmark and foundation model for complicated chart reasoning. *arXiv preprint arXiv:2402.12185*, 2024. 2, 5, 6

[48] Ximing Xing, Juncheng Hu, Guotao Liang, Jing Zhang, Dong Xu, and Qian Yu. Empowering llms to understand and generate complex vector graphics. *arXiv preprint arXiv:2412.11102*, 2024. 3

[49] Zhengzhuo Xu, Sinan Du, Yiyan Qi, Chengjin Xu, Chun Yuan, and Jian Guo. Chartbench: A benchmark for complex visual reasoning in charts. *arXiv preprint arXiv:2312.15915*, 2023. 2, 5, 6, 7, 14, 17

[50] Zhengzhuo Xu, Bowen Qu, Yiyan Qi, Sinan Du, Chengjin Xu, Chun Yuan, and Jian Guo. Chartmoe: Mixture of expert

connector for advanced chart understanding. *arXiv preprint arXiv:2409.03277*, 2024. 2

[51] Xudong Yang, Yifan Wu, Yizhang Zhu, Nan Tang, and Yuyu Luo. Askchart: Universal chart understanding through textual enhancement. *arXiv preprint arXiv:2412.19146*, 2024. 2

[52] Zhiyu Yang, Zihan Zhou, Shuo Wang, Xin Cong, Xu Han, Yukun Yan, Zhenghao Liu, Zhixing Tan, Pengyuan Liu, Dong Yu, et al. Matplotagent: Method and evaluation for llm-based agentic scientific data visualization. *arXiv preprint arXiv:2402.11453*, 2024. 3

[53] Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li, Weilin Zhao, Zhihui He, et al. Minicpm-v: A gpt-4v level mllm on your phone. *arXiv preprint arXiv:2408.01800*, 2024. 6, 7

[54] Jiabo Ye, Anwen Hu, Haiyang Xu, Qinghao Ye, Ming Yan, Guohai Xu, Chenliang Li, Junfeng Tian, Qi Qian, Ji Zhang, et al. Ureader: Universal ocr-free visually-situated language understanding with multimodal large language model. *arXiv preprint arXiv:2310.05126*, 2023. 2

[55] Zhenfei Yin, Jiong Wang, Jianjian Cao, Zhelun Shi, Dingning Liu, Mukai Li, Xiaoshui Huang, Zhiyong Wang, Lu Sheng, Lei Bai, et al. Lamm: Language-assisted multi-modal instruction-tuning dataset, framework, and benchmark. *Advances in Neural Information Processing Systems (NeurIPS)*, 36, 2024. 2

[56] Fatemeh Pesaran Zadeh, Juyeon Kim, Jin-Hwa Kim, and Gunhee Kim. Text2chart31: Instruction tuning for chart generation with automatic feedback. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 11459–11480, 2024. 2

[57] Xingchen Zeng, Haichuan Lin, Yilin Ye, and Wei Zeng. Advancing multimodal large language models in chart question answering with visualization-referenced instruction tuning. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 2024. 2

[58] Jieyu Zhang, Le Xue, Linxin Song, Jun Wang, Weikai Huang, Manli Shu, An Yan, Zixian Ma, Juan Carlos Niebles, Caiming Xiong, et al. Provision: Programmatically scaling vision-centric instruction data for multimodal language models. *arXiv preprint arXiv:2412.07012*, 2024. 2

[59] 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. *arXiv preprint arXiv:2404.16635*, 2024. 2

ECDBench tests a different chart mix, with human-checked questions.

Effective Training Data Synthesis for Improving MLLM Chart Understanding

Supplementary Material

A. Chart Types and Subplot Combinations

As illustrated in Fig. 5, our ECD dataset includes 29 distinct chart types for single plot, *i.e.,* "line", "bar", "pie", "area", "errorpoint", "treemap", "funnel", "node", "density", "histogram", "box", "bubble", "candlestick", "heatmap", "radar", "rose", "3d", "errorbar", "quiver", "scatter", "violin", "contour", "bar + line overlay", "pie + bar overlay", "histogram + density overlay", "violin + box overlay", "scatter + histogram overlay", "scatter + density overlay" and "hexbin + hist overlay".

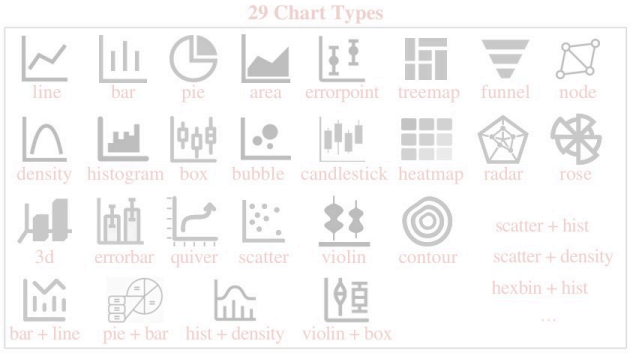


Figure 5. Distribution of Chart Types in our proposed effective chart dataset (ECD). Our dataset encompasses 29 diverse chart types, surpassing the variety present in existing datasets.

Additionally, for the combinations of subplots, our dataset only uses random combinations of 2 single-plot chart types (excluding overlays, can be the same type, *e.g.,* "line + line"). Following the pipeline outlined in this paper, it can be seamlessly extended to incorporate combinations of more chart types in practice, such as combinations involving 3 chart types, (*e.g.,* "line + bar + pie", *etc.*). Ultimately, 252 distinct combinations are successfully generated (excluding unsuccessful examples after filtering).

B. Chart Themes and Layout Distribution

As shown in Fig. 6, we present the distribution of chart themes and layouts in the proposed ECD dataset.

Chart Themes Distribution. Our dataset contains 25 scientific themes, with an average of approximately 421.4 images per theme. The number of images in each theme is as follows: Art and Design: 450, Agriculture: 453, Geography: 359, Medicine: 423, Engineering: 485, Law: 383, Biology: 451, Sports: 399, Mathematics: 495, Environmental Science: 450, Anthropology: 406, Sociology: 379, Computer Science: 474, Education: 380, Architecture: 405, Psychology: 436, Economics: 345, Statistics: 415, History: 332, Chemistry: 501, Media and Journalism: 383, Finance: 412, Physics: 500, Astronomy: 414, Linguistics: 405. Among them, the Chemistry theme has the highest count, while the History theme has the lowest.

Chart Layouts Distribution. We include 13 primary layouts, ranging from "1 by 1 (1, 1)" to "4 by 2 (4, 2)". The "1 by 1" layout represents a single chart image, with a total of 6,414 instances (accounting for 60.9%), while the other layouts represent multiple images, comprising 4,121 instances (accounting for 39.1%). This includes 631 instances of multiple images of the same chart type (*e.g.,* line + line subplots) and 3,490 instances of multiple images of different chart types (*e.g.,* line + bar subplots).

Models	ECDBench		
	Reasoning	Descriptive	Average
Claude Series			
Claude 3.5 Sonnet [5]	41.99	68.14	55.07
Claude 3.7 Sonnet [6]	43.38	**69.61**	56.50
Claude 4 Sonnet [7]	**44.20**	69.36	**56.78**
Gemini Series			
Gemini-2.5-Pro [13]	**44.36**	**76.88**	**60.62**
GPT Series			
Random (GPT-4o) [4]	4.58	1.63	3.10
GPT-4o-mini [4]	24.26	57.27	40.77
GPT-4o [4]	35.62	70.18	52.90
o1 [20]	40.52	74.18	57.35
o3 [40]	56.13	74.51	65.32
o4-mini [40]	**57.03**	**77.45**	**67.24**
Qwen-VL Series			
Qwen2.5-VL-7B [8]	19.04	57.35	38.19
Qwen2.5-VL-32B [8]	24.92	53.92	39.42
Qwen2.5-VL-72B [8]	**38.81**	**68.46**	**53.64**

Table 8. Performance comparison of several closed-source and open-source MLLMs on ECDBench. Among the listed models, o4-mini achieved the best performance.

C. ECDBench Statistics and More Results

The ECDBench we constructed comprises a total of 1,224 chart images, including 364 single-plot charts and 860 combined subplot charts. Among the combined-subplot charts, 457 feature combinations of two chart types (*e.g.,* "bar + line". with various layouts ranging from (1, 2) to (3, 3) / (4, 2), *etc.*), while 403 involve combinations of three chart types (*e.g.,* "bar + pie + scatter". with (1, 3) or (3, 1) layout). The average size of images in ECDBench is 1378 × 968 px. ECDBench includes 2,448 QA pairs, with each image ac-

The benchmark has 364 single-plot and 860 combined-plot charts.

1,224 charts → descriptive / reasoning } 2,448 QA pairs

Cheap questions do not guarantee transfer beyond chart understanding.

companied by 1 descriptive QA and 1 reasoning-based QA. For the manual checks and annotations of ECDBench, we employ PhD students who are familiar with common scientific charts.

Tab. 8 presents a more detailed comparison of evaluation results on several state-of-the-art (SOTA) proprietary and open-source MLLMs, including the Claude series, Gemini series, GPT series, and Qwen-VL series. Among all the evaluated MLLMs, o4-mini consistently achieves the highest performance across all three metrics (57.03% reasoning acc, 77.45% descriptive acc and 67.24% average acc).

D. FID and Avearage Pixel Entropy

In this paper, we employ the FID (Fréchet Inception Distance) and Average Pixel Entropy as metrics to assess the image realism and complexity of chart datasets respectively. Here, we present the specific calculation formula.

Fréchet Inception Distance (FID): For the real dataset I_{real} and Synthetic dataset I_{syn}:

$$\text{FID}(I_{\text{real}}, I_{\text{syn}}) = \|\mu_{\text{real}} - \mu_{\text{syn}}\|^2 + \text{Tr}\left(\Sigma_{\text{real}} + \Sigma_{\text{syn}} - 2\sqrt{\Sigma_{\text{real}}\Sigma_{\text{syn}}}\right), \tag{1}$$

where μ_{real} and μ_{syn} are the mean feature vectors of the Inception network embeddings, Σ_{real} and Σ_{syn} are their covariance matrices, and $\text{Tr}(\cdot)$ denotes the matrix trace operator.

Average Pixel Entropy: The average pixel entropy can be defined as:

$$H(x) = -\sum_{i=1}^{n} p(i) \log p(i), \tag{2}$$

$$\text{AvgEnt}(I) = \frac{1}{N} \sum_{k=1}^{N} H(x_k), \tag{3}$$

where $p(i)$ is the normalized histogram probability of intensity value i, n is the number of discrete intensity levels (*e.g.,* 256 for 8-bit images), $H(x_k)$ is entropy of the k-th image.

E. Computational and GPT-4o Cost

Dataset	#Images	#QA Pairs	$Per-image	$Per-QA	$Total
CharXiv	2.3k	5k	0.37	0.170	∼0.85k
ReachQA	3.2k	19.9k	0.10	0.016	∼0.31k
ECD	10.5k	321.5k	0.20	0.007	∼2.15k

Table 9. Cost ($) comparison between ECD and existing datasets.

Generating 1,000 single-plot images takes an average of 7.1 seconds per image and 22.40 USD in total. On the other hand, generating 1,000 combined subplots costs 22.6 seconds per image and 85.12 USD in total. During diversification, each set of 1,000 images costs 30.14 USD and requires an average of 40.18 seconds per image. The figure size post-processing and the image filtering costs 27.90 USD per 1,000 images in total. Additionally, generating QA pairs for 1,000 chart images costs 38.97 USD. In total, generating ECD costs ∼2,145 USD. We add cost comparison in Tab. 9. While ECD has the highest total cost due to its large number of image and QA pairs, its per QA cost is the lowest, and its per-image cost is medium (but not rediculously high on 0.2$ per image), because it has many multi-subplot charts, which is essential in its design. We also tried replacing GPT-4o with Qwen2.5-VL-32B for training data generation. Results are shown in Tab. 10. We find that both yield improvement on ReachQA benchmark, but the dataset generated by Qwen2.5-VL-32B does not improve performance on CharXiv.

The reported synthesis cost is about $0.007 per retained QA pair.

Models	CharXiv	ReachQA
LLaVA-Next-Llama3-8B	35.06	15.65
+ ECD by Qwen2.5-VL-32B	34.90	21.40
+ ECD by GPT-4o	45.66	18.85

generator choice matters

Table 10. Comparison on ECD generation (5k images, 30k QAs).

F. Cross-task MLLM SFT Impact

To assess the impact of our synthetic chart QA data on model generalization, we fine-tune LLaVA-Next-Llama3-8B on our dataset and evaluate on MathVista [33] (math images, testmini split), MMBench [31] (mixed natural and scientific images, dev-en split), and RealworldQA [45] (natural images, test split). In Tab. 11, our model leads to slight improvements on MathVista and RealworldQA and a decrease on MMBench. These early results might be inconclusive, and we plan to expand them in the future.

Models	Mathvista [33]	MMBench [31]	RealWorldQA [45]
LLaVA-Next-Llama3-8B	37.00	**79.03**	58.95
+ ECD	**38.40**	77.36	**59.22**

Table 11. GPT-Acc (%) on 3 common vision-language benchmarks.

one cross-task decline

G. Analysis of the Generation Pipeline

We construct several variants to demonstrate the effectiveness of the proposed pipeline.

Separate data and code generation vs. joint generation is shown in Tab. 12 (a) vs Tab. 12 (c). We separately generate the code and data, comparing it with the case where the code and data are generated together. We observe that ours is superior to generating data and code jointly, evidenced by the lower FID and higher entropy.

Visual similarity is useful evidence, but its measurement has a domain gap.

Generation Strategy	FID↓	Avg. Entropy↑
(a) Code and Data Joint Gen.	126.86 ± 3.11	1.91
(b) Subplots Parallel Gen.	123.26 ± 2.15	1.90
(c) Ours	**120.63 ± 3.23**	**2.05**

Table 12. Comparison of three chart generation strategies. (a) Generating chart code and data simultaneously. (b) Generate all the subplots data simultaneously. (c) Our data generation method (decoupling code and data; conditional generation of subplots). We define 12 different layouts for the common "bar+line" combination, and generate 10 images for each layout. We use FID (with CharXiv) [42] and average pixel entropy to evaluate the realism and complexity of generated chart images.

Effectiveness of conditional and sequential data generation for multi-subplot charts. We condition each subplot on previous ones to maintain consistency. We compare this method with joint multi-subplot generation in Tab. 12 (b) vs Tab. 12 (c). We observe a decline in the realism and complexity if we do joint generation, which indicates the effectiveness of our design.

H. Numerical Diversity, Chart Style Statistics and Multi-Subfigure Reasoning Analysis

Numerical Diversity. Our current method has three strategies for numerical diversity. 1) We randomly sample data trend, *e.g.*, increasing, decreasing, or stable. 2) we randomly sample the number of elements, *e.g.*, the number of groups in a bar plot. 3) We use a high temperature (1.0) in GPT-4o to encourage varied parameter output.

Chart Style Statistics. We calculate the percentage of charts undergoing various changes using our synthesis method. For line-charts we find: 100.00% charts have changes in line color, 99.65% in font size, 95.09% in line width, 67.37% with non-default background, 59.30% with transparency, 47.37% in grid, 39.65% in new annotations, 29.12% in area shading, 17.19% include arrows, 11.23% draw threshold lines, 10.18% remove axis borders, 10.18% add zoom-in insets and 3.51% include error bars. In total, 58.95% of line charts include 6 or more distinct changes.

Only about 13.4% of 500 sampled reasoning questions were judged to need multiple subfigures.

Multi-Subfigure Reasoning QA Pairs. We analyze 500 reasoning QA pairs. ∼13.4% are identified by GPT-4o as requiring multi-subfigure reasoning. Moreover, while some QAs can be answered by a single figure in a multi-subfigure chart, such QAs are still difficult because of the need for subfigure localization.

I. Discussions and Limitations

What makes a high-quality chart training set? This paper mainly uses FID and average entropy as indicators to the similarity to real images and image complexity. However, the fact that ECD performs better than other datasets does not mean that ECD always has a lower FID and higher average entropy. We validated that the number of themes and chart types are also useful. Other factors such as text formats, font sizes, and even colors may also influence the training set quality.

Lack of a dedicated vision encoder for chart. FID calculation requires feature extraction using InceptionNet-V3, pretrained on ImageNet (a natural image dataset). There are domain gaps between natural images and chart images, which might limit the efficacy of FID. The community still lacks a proper chart vision encoder to better compute indicators that require feature extraction.

J. Comparison with Other Chart Datasets

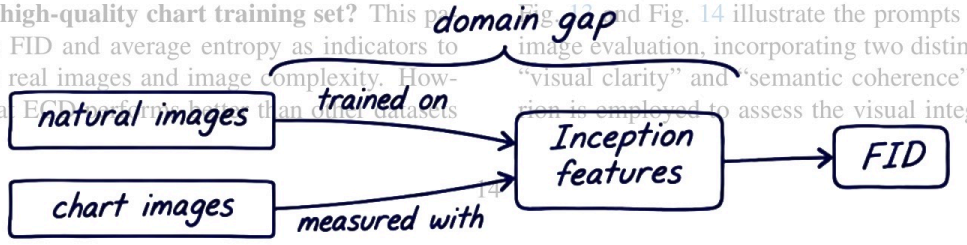
To facilitate further comparison with other datasets, Fig. 7 present qualitative comparisons. This figure presents a visual comparison between our dataset's chart images and those from several chart QA training datasets, including PlotQA [39], ChartQA [34], ChartBench [49], SimChart9k [46], ChartAssistant [38], and ReachQA [15]. Compared to these datasets, ours features a greater variety of chart types, more diverse combinations, more intricate details, and richer visual representations.

K. Prompt for Chart Image Generation

Fig. 8, Fig. 9, Fig. 10 and Fig. 11 illustrate the prompts we use to guide GPT-4o in generating single plot chart, combined subplot charts, and diverse visual variations. I) *Single plot chart*: We employ predefined chart functions and examples to prompt GPT-4o to generate the data for a single plot chart. II) *Combined subplot chart*: each single plot in the combined chart is generated sequentially based on conditional inputs (*i.e.*, data of previous generated plots), ensuring semantic continuity across the entire chart composition. III) *Diversification*: We define 5 and 6 diversification strategies for the single-plot chart and the combined subplot chart, respectively, and randomly select 1 strategy for each chart generation. It is important to note that additional prompts are incorporated to guide GPT-4o in prioritizing detail modifications while ensuring the preservation of the originally generated data. Fig. 12 presents the post-processing prompt applied to chart images, designed to empower GPT-4o in refining figure size and resolution. This step aims to enhance the conveyance of visual information, particularly in improving the clarity of elements such as text.

L. Prompt for Chart Image Rating

Fig. 13 and Fig. 14 illustrate the prompts utilized for chart image evaluation, incorporating two distinct rating criteria: "visual clarity" and "semantic coherence". The first criterion is employed to assess the visual integrity of the chart



(on a scale from 1 to 5), specifically targeting the exclusion of images exhibiting excessive blank spaces or chaotic overlaps. The second criterion evaluates the semantic coherence (also on a scale of 1 to 5), ensuring that the chart adheres to a well-defined and logically structured theme. The integration of these two metrics serves to filter out charts of insufficient quality.

M. Prompt for Chart QA Generation

Fig. 15 and Fig. 16 illustrate the prompts we provide to GPT-4o for generating question-answer (QA) pairs based on chart code, where chart images are also incorporated as input to assist in determining whether the questions can be answered solely based on visual information. The generated QA pairs are categorized into two types: descriptive and reasoning. Descriptive QA pairs focus on the identification of fundamental textual, numerical, and graphical elements within the chart. Reasoning QA pairs, on the other hand, emphasize hierarchical analysis, computational inference, and logical deduction, with the generated answers incorporating rationale as part of the reasoning process. Additionally, for each generated QA pair, GPT-4o is tasked with producing a confidence rating, which serves as a measure for subsequent filtering.

N. Prompt for Evaluation on ECDBench

Fig. 17 illustrates the prompt used for evaluating model predictions on ECDBench. Similarly, we employ the LLM-as-Judge approach to assess the responses. For answers containing numerical results, a 5% tolerance is permitted to account for visual recognition discrepancies.

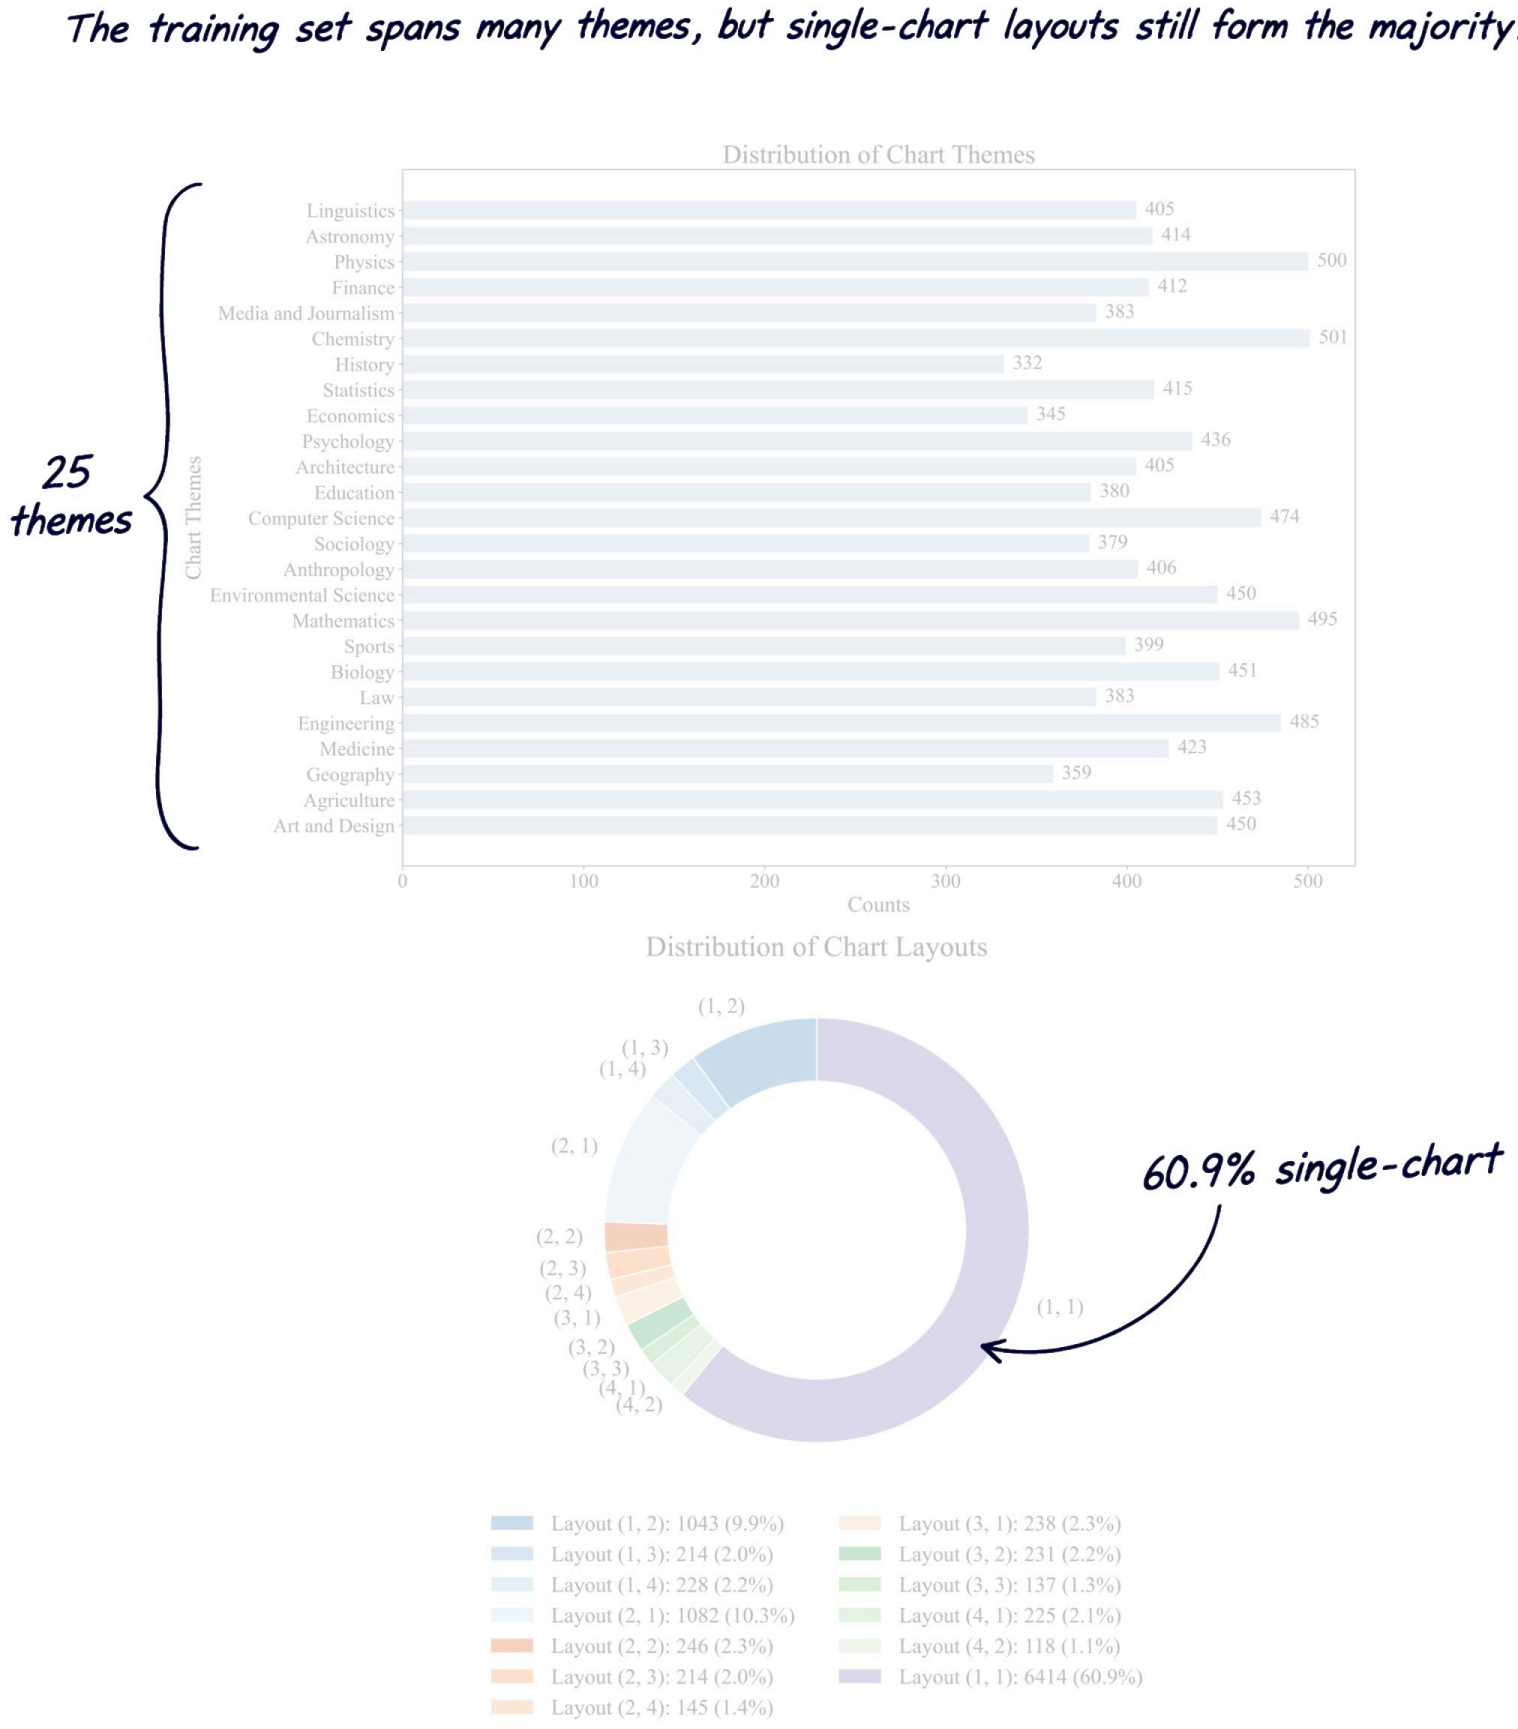


Figure 6. Distribution of themes and layouts in our proposed effective chart dataset (ECD). Our dataset comprises 25 distinct themes and 12 multi-plot layouts. It includes 6,414 single-chart images and 4,121 combined subplot images. Among the multi-plot layouts, 631 subplots consist of combinations of the same chart type, while 3,490 subplots involve combinations of different chart types. Following the pipeline proposed in this paper, our dataset can be further expanded to include more chart types and layouts.

The examples show a broader visual vocabulary, rather than a single chart template.



Figure 7. Compared to the chart images in existing chart QA training datasets (sampled for better visualization, *i.e.*, PlotQA [39] (3 types), ChartQA [34] (3 types), ChartBench [49] (9 types), SimChart9k [46] (3 types), ChartAssistant [38] (9 types), ReachQA [15] (10 types)), it is evident that our ECD (29 types) exhibits a greater diversity of subplot combinations (>250 types of combination) and complexity.

Illustrative samples show variety; accuracy still needs separate evaluation.

Prompt for Single Plot Generation

Assistant:
Intelligent Assistant Prompt: Input Synthesis Helper

You are a "Function Input Synthesis Helper," and your task is to generate an appropriate function call based on a provided function, filling the call with reasonable parameters.

User-Provided Function:
```python
{user_function}
```

Synthesized Function Call Example:
```python
{function_name}(
    {parameters_and_values}
)
```

Important Considerations:
--Ensure that data_list contains realistic data with variability and noise.
--The pattern for parameters should reflect potential trends, cycles, randomness, and correlations.
--Data should exhibit complexity and noise, avoiding simple linear or polynomial.
--Make sure the elements in data_list match labels.
--Parameters such as colors_list, linestyles_list, and markers_list should match the structure and amount of data in data_list.
--Each parameter should follow a unique potential pattern.
--Users may provide data themes and constraints on certain parameters, and data should be generated accordingly.
--Data values should make sense and fall within appropriate ranges for the given context.
--Labels should be specific and meaningful, avoiding generic terms like "Product One" or "Product Two." The names and chart labels should be specific to the data and context.
--Chart should be clear and visually appealing, with appropriate colors and styles for each line chart. linewidths should larger than 1 for better visibility.
--Adjust the width of any elements (e.g., bars, lines) to ensure they are visible and distinguishable.
--Generate **random and diverse** colors/styles/markers, and don't start with **a fixed color/style/marker**. Explore elegant and unique color tones suitable for research papers, while also considering commonality.
--You MUST add noise to the data to make it more realistic and challenging.
--Do NOT generate linear data.
--DO NOT return function code, just the function call with parameters.

User:
{plot_function}
Example Usage:
{example_usage}
Chart Theme: Imagine a real application scenario under the **{random_theme}** theme and generate the data accordingly. Be specific. For example, if the theme is Finance, you could imagine a company struggling financially but expanding into new business areas. If the theme is Astronomy, you could picture a star nearing the end of its life.
The number of elements:
First level: {random.randint(3, 6)};
If second level make sense in the context, Second level: {random.randint(3, 6)}.
Trend of data: create {trend_string}.
grid: {random.choice([False, True])}.
save_path: {png_save_path}.

Theme Collections:
["Economics", "Psychology", "Sociology", "Biology", "Education", "Engineering", "Law", "Astronomy", "Computer_Science", "Geography", "Physics", "Chemistry", "History", "Environmental_Science", "Anthropology", "Media_and_Journalism", "Mathematics", "Statistics", "Finance", "Medicine", "Art_and_Design", "Agriculture", "Linguistics", "Architecture", "Sports"]

Figure 8. Prompt template for generating a single plot. We define 25 academic themes for generating charts across diverse disciplines.

Prompt for Single Plot Diversification

Assistant 1:
Intelligent Assistant Prompt: Single Plot Code Synthesis and Diversification Helper

You are a 'Single Plot Code Synthesis and Diversification Assistant.' Your task is to synthesize an appropriate single `plot` code based on the predefined chart function and calling example provided by the user, ensuring that it follows the user's diversification strategy to make the synthesized code sufficiently varied.

User-Provided Plot Function:
```python
{single_plot_function}
```

User-Provided Calling Example:
```python
{calling_example}
```

Synthesized Code Example:
```python
{code_body}
```

Assistant 2:
Intelligent Assistant Prompt: Code Debug Helper

You are a 'Code Debug Assistant.' Your task is to identify and fix issues in the user's code based on any provided errors, ensuring it works correctly.

User-Provided Code:
```python
{user_code}
```

Error Message:
{error_message}

Returned Code:
```python
{whole_code_after_fix}
```

User:
{single_plot_function}
Provided calling example for single plot: {calling_example}. The data/labels/legends in the chart must be preserved at the diversification stage.
Diversification Strategy for Single Plot Chart Code: {random.choice{Diversification_Strategy_Collections}}.
save_path: {div_png_save_path}.
Important Considerations:
--Code Diversity: Provide multiple styles and techniques for modifying the chart, using different libraries such as Seaborn to enhance the code and visual presentation. Must avoid returning results as **functions like 'def' or function calls**, and instead, provide complete executable code.
--Layout and Organization: You must follow the **save_path and diversification strategy** and ensure clear separation (without any overlap), alignment and labeling for readability. Additionally, feel free to adjust the figsize to ensure that all elements are fully visible and can be displayed.
--Original Data Preservation: You are required to create the code, without modifying or ignoring the original chart data. It is worth noting that for the line_num and bar_num charts, the number annotations should not be modified or removed.
Note: Be sure not to use any interactive elements that cannot be saved. If chart type is Funnel, use plotly lib for implementation.

Diversification_Strategy_Collections:
["Add text labels, annotations, arrows, uncertainty bars, threshold lines (which should have specific names, not just 'threshold'; these could include curves like exponential/logarithmic curves, and not always be straight lines), or highlights (e.g., using a circle or highlighting a range in a specific color) to emphasize key data points, trends, or regions, ensuring that these annotations are contextually relevant and not generic. The number of items added can exceed one.",
"Modify the font styles, colors or sizes significantly for titles, labels or ticks.",
"Use gradient fills, area shading (typically along the line itself, within a defined range above and below it, must not between the line and the x-axis), or transparency effects to enhance depth. Additionally, fine-tune grid lines, background colors, and shadow effects to improve visual appeal.",
"(If Applicable) Remove axis borders for a cleaner, modern look.",
"(If Applicable) Incorporate a **zoomed-in inset** of a particular section, ensuring the area is appropriately sized (usually a very small size) and placed, making sure that the elements are visually separated without overlapping."]

Figure 9. Prompt template for single plot diversification. We define 5 diversification strategies and randomly select one for each plot.

Prompt for Combined Subplots Generation and Diversification – (a)

Assistant 1:
Intelligent Assistant Prompt: Input Synthesis Helper

You are a "Function Input Synthesis Helper," and your task is to generate an appropriate function call based on a provided function, filling the call with reasonable parameters.

User-Provided Function:
```python
{user_function}
```

Synthesized Function Call Example:
```python
{function_name}(
    {parameters_and_values}
)
```

Important Considerations:
--Ensure that data_list contains realistic data with variability and noise.
--The pattern for parameters should reflect potential trends, cycles, randomness, and correlations.
--Data should exhibit complexity and noise, avoiding simple linear or polynomial.
--Make sure the elements in data_list match labels.
--Parameters such as colors_list, linestyles_list, and markers_list should match the structure and amount of data in data_list.
--Each parameter should follow a unique potential pattern.
--Users may provide data themes and constraints on certain parameters, and data should be generated accordingly.
--Data values should make sense and fall within appropriate ranges for the given context.
--Labels should be specific and meaningful, avoiding generic terms like "Product One" or "Product Two." The names and chart labels should be specific to the data and context.
--Chart should be clear and visually appealing, with appropriate colors and styles for each line chart. linewidths should larger than 1 for better visibility.
--Adjust the width of any elements (e.g., bars, lines) to ensure they are visible and distinguishable.
--Generate **random and diverse** colors/styles/markers, and don't start with **a fixed color/style/marker**. Explore elegant and unique color tones suitable for research papers, while also considering commonality.
--You MUST add noise to the data to make it more realistic and challenging.
--DO NOT generate linear data.
--DO NOT return function code, just the function call with parameters.

Assistant 2:
Intelligent Assistant Prompt: Combined Subplot Chart Code Synthesis Helper

You are a 'Combined Subplot Chart Code Synthesis Assistant.' Your task is to synthesize an appropriate combined subplot chart code based on the individual chart functions and calling examples provided by the user, ensuring that all functions and examples integrate seamlessly.

User-Provided Functions:
```python
{user_function}
```

User-Provided Calling Examples:
```python
{user_calling_examples}
```

Synthesized Code Example:
```python
{code_body}
```

Assistant 3:
Intelligent Assistant Prompt: Code Debug Helper

You are a 'Code Debug Assistant.' Your task is to identify and fix issues in the user's code based on any provided errors, ensuring it works correctly.

User-Provided Code:
```python
{user_code}
```

Error Message:
{error_message}

Returned Code:
```python
{whole_code_after_fix}
```

Figure 10. Prompt template for combined subplots generation and diversification - Part (a).

Prompt for Combined Subplots Generation and Diversification – (b)

==========Stage I：Synthesize the data of all individual Subplots==========

User:
{current_plot_function}
Example Usage of single chart: {current_example}.
Subplot Data generated previously: {before_data}.
Color Scheme: Ensure that the color scheme across subplots is consistent, balancing both aesthetic appeal and clarity in data presentation.
Chart Theme: Imagine a real application scenario under the **{random_theme}** theme and generate the data accordingly. Be specific. For example, if the theme is Finance, you could imagine a company struggling financially but expanding into new business areas. If the theme is Astronomy, you could picture a star nearing the end of its life.
The number of elements:
First level: {random.randint(3, 6)};
If second level make sense in the context, Second level: {random.randint(3, 6)}.
Trend of data: create {trend_string}.
grid: {random.choice([False, True])}.
save_path: {png_save_path}.
Important Considerations: You must generate parallel sub-graphs that maintain strong internal cohesion while introducing meaningful variations, ensuring both similarity and distinctiveness within the data structure.

==========Stage II：Combine Subplots and Diversification==========

User:
{all_plot_functions}
Provided chart data for each subplot: {combined_data}. The data/labels/legends in the chart must be preserved at the combining stage.
Layout: A {random_layout[0]}x{random_layout[1]} grid of subplots, totaling {random_layout[0] * random_layout[1]} subplots. You must ensure the elements in the data_list match the number of subplots in the layout.
save_path: {combined_png_save_path}.
Diversification Strategy for Combined Subplots Chart Code: {random.choice{Diversification_Strategy_Collections}}.
Important Considerations:
--Code Diversity: Provide multiple styles and techniques for modifying the chart, using different libraries such as Seaborn to enhance the code and visual presentation. Must avoid returning results as **functions like 'def' or function calls,** and instead, provide complete executable code.
--Layout and Organization: You must follow the **requested layout, save_path and change strategy** and ensure clear separation (without any overlap), alignment and labeling for readability. Additionally, feel free to adjust the figsize to ensure that all elements are fully visible and can be displayed.
--Original Data Preservation: You are required to create the code, without modifying or ignoring the original chart data. It is worth noting that for the line_num and bar_num charts, the number annotations should not be modified or removed.
Note: Be sure not to use any interactive elements that cannot be saved. Besides, {random.choice{Subtitles_Strategy_Collections}}.

Diversification_Strategy_Collections:
["Add the title of the whole combined chart in a specific and meaningful way, avoiding general terms.",
"Add text labels, annotations, arrows, uncertainty bars, threshold lines (which should have specific names, not just 'threshold'; these could include curves like exponential/logarithmic curves, and not always be straight lines), or highlights (e.g., using a circle or highlighting a range in a specific color) to emphasize key data points, trends, or regions, ensuring that these annotations are contextually relevant and not generic. The number of items added can exceed one.",
"Modify the font styles, colors or sizes significantly for titles, labels or ticks.",
"Use gradient fills, area shading (typically along the line itself, within a defined range above and below it, must not between the line and the x-axis), or transparency effects to enhance depth. Additionally, fine-tune grid lines, background colors, and shadow effects to improve visual appeal.",
"(If Applicable) Remove axis borders for a cleaner, modern look.",
"(If Applicable) Incorporate a **zoomed-in inset** of a particular section, ensuring the area is appropriately sized (usually a very small size) and placed, making sure that the elements are visually separated without overlapping."]

Subtitles_Strategy_Collections:
["leave the subfigure titles as they are without any letters",
"retain the original subfigure titles and add letters before the titles, such as (a), (b), (c), (d)",
"choose to remove the original subfigure titles and use only the letters"]

Figure 11. Prompt template for combined subplots generation and diversification - Part (b). The combined subplots introduce an additional diversification strategy on top of the single plot (i.e., add the title of the whole chart). Additionally, the diversification of subtitles is also randomly selected from the collection.

Prompt for Figure Size Post-Processing

You are an experienced Python developer specializing in data visualization and chart rendering. The following Python code generates and saves charts. Some charts may have unnecessarily large or inappropriate sizes, while others may already be optimal. Modify the code as follows:

1. If the chart size is inappropriate (e.g., unnecessarily large or excessively small), adjust the figsize to a **reasonable resolution** suitable for typical use cases.
2. If the chart size is already appropriate, leave it unchanged.
3. Saves with dpi **no higher than 150** (e.g., `dpi=100`).
4. Guarantee that axis labels, titles, legends and tick labels **do not overlap**:
   - Use `fig.tight_layout()` or `constrained_layout=True`.
   - Dynamically scale `figsize` or reduce `plt.rcParams["font.size"]` when needed.
   - Verify that every text element remains readable and non-overlapping.
5. Do not alter the chart content, style, or any other functionality of the code.
6. Headless: set backend "Agg" before pyplot and remove all **.show()** calls.
7. Ensure the modified code is efficient, readable, and adheres to best practices for saving visualization outputs.

Here is the original code:
```python
{original_code}
```
Respond with the complete modified code, enclosed in Python code block tags for clarity.

Figure 12. Prompt template for figure size post-processing. We enable GPT-4o to dynamically adjust figure size and visual elements during post-processing, ensuring better clarity and visualization.

Prompt for Visual Clarity Rating

<Chart_Image>

You're an expert evaluating the **visual clarity** of a chart. Assign a rating from 1 to 5, adhering to the detailed descriptions below:

1 point: The chart is essentially unreadable due to overwhelming issues, such as excessive clutter, severe text overlap, missing legends, or large blank areas, making it impossible to interpret the data accurately.
2 points: The chart has significant design flaws that hinder understanding. Issues like noticeable clutter, text overlap, or missing elements complicate the data's presentation and make it difficult to extract meaningful insights.
3 points: While the chart communicates the data to some extent, it suffers from moderate issues. Mild clutter, misleading scales, or slight text overlap may confuse viewers, affecting overall clarity.
4 points: The chart is largely clear and conveys the data effectively, though minor issues like slight text overlap, small misalignments, or minor clutter might cause slight distractions without greatly affecting comprehension.
5 points: The chart is visually excellent, with precise and well-organized data presentation. It uses space effectively, ensuring that the information is easy to interpret and completely free of any errors.

Layout for current chart: **{layout}**, you should consider whether the chart image matches the provided layout. Give your rating for the provided chart image now.

Provide your rating in the format without any explanation:
Rating: (integer value within 1 to 5)

Figure 13. Prompt template for visual clarity rating. We apply a 1∼5 point scale for GPT-4o to evaluate and select chart images with excellent visual clarity. The underlined portions are applicable only to combined subplot charts with layout that are not 1 by 1.

Prompt for Semantic Coherence Rating

<Chart_Image>

You're an expert evaluating the **semantic coherence** of a chart. Assign a rating from 1 to 5, adhering to the detailed descriptions below:

1 point: The chart is irrelevant to its intended message or theme. The data is disjointed or disconnected, with no clear connection between elements or subplots. There is no recognizable story, and the chart does not support the overall narrative in any meaningful way.
2 points: The chart's content partially aligns with its theme, but there are significant gaps in how the data is presented. The relationships between elements are weak or unclear, and subplots may seem out of place, making it difficult to extract a clear narrative or insight from the chart.
3 points: The chart generally supports its theme or message, but the connections between data points or subplots are inconsistent or ambiguous. The overall story is somewhat discernible, but there are sections where the data feels disjointed or hard to follow.
4 points: The chart aligns well with its theme, presenting a mostly coherent story. There are clear connections between the main plot and subplots, though some minor ambiguities or distractions might slightly affect the clarity of the narrative. The data supports the theme effectively.
5 points: The chart is highly coherent, with a clear, logical progression of ideas. The data is directly relevant to the theme, and the relationships between the main plot and any subplots are seamless. The chart presents a compelling and easy-to-follow narrative, making the insights immediately apparent.

Theme for current chart: **{theme}**, you should consider whether the chart image matches the provided theme. Give your rating for the provided chart image now.

Provide your rating in the format without any explanation:
Rating: (integer value within 1 to 5)

Figure 14. Prompt template for semantic coherence rating. We apply a 1∼5 point scale for GPT-4o to evaluate and select chart images with clear, logical semantic coherence.

Prompt for Descriptive QA Pairs Generation

User:
You are a careful and precise observer of the chart. You are provided with the original chart code as follows:

```python
{chart_code}
```

Your task is to carefully analyze the chart code details and generate **Basic Descriptive-based** question-answer pairs. Additionally, a chart image is provided to assist in verifying whether the questions can be fully addressed based on the visual information it contains.

Areas to Focus On:
1. **Textual Content**: This includes elements like titles, subtitles, axis label (e.g., x-axis or y-axis), tick marks (e.g., labeled ticks at extremes, like **the leftmost/rightmost/highest/lowest**, etc.), annotations, or legends (e.g., list of legend labels).

2. **Numerical Content**: This refers to the count of variables (e.g., lines/bars/instances), the number of ticks across all axes, **differences between numerical tick values on the x- and y-axes**, label count in legends, and the maximum/minimum (or **difference between the maximum and minimum**) and intervals in continuous colorbar, etc.

3. **Graphical Content**: This includes visual characteristics such as:
   - Chart type
   - Point positions
   - Line intersections and thresholds
   - Overall orientation (e.g, "veritical" or "horizontal", not the angle of specific elements)
   - **Subplot layouts** (e.g., answer in format like "1 by 2", "2 by 4", "3 by 3")
   - Number of subplots
   - Highlights and zoom-in inset (e.g., emphasizing specific features for clarity)
   - General trends and patterns (e.g., "increase," "decrease," "stabilize")

Response Format:
Please use the following JSON format to ensure clarity:
```json
{
   {
     "image_id": "{img_id}_1",
      "rating": "Rate your confidence in this question and answer [0-5], Integer",
      "question": "Provide your question here",
      "answer": "Provide your corresponding complete answer here"
    },
    {
      "image_id": "{img_id}_2",
        ...
     },
        …
}
```

Notes:
1. Focus only on aspects that can be inferred visually from the chart image. Do not refer to facts or terms that directly suggest the content is based on the code rather than the chart image itself.
2. The total number of questions should be **10 to 15,** and feel free to rephrase your questions to enhance diversity.
3. The QA pairs should include questions involving the enumeration and counting of elements such as legends, ticks (e.g., total across all axes, like x_axis + y_axis), variables, etc. Furthermore, simple calculations are necessary for identifying differences.
4. Do not include questions about **alpha value, transparency, opacity, padding, spacing, grid, width, viewing angle, rotation, function (e.g.,'np.random'), font size, figsize, save filename and save_path** in the code. Focus only on the meaningful visual elements in the chart and avoid mentioning 'code' in the questions and answers.
5. Ensure the image_id numbers increment sequentially, such as _1, _2, _3, _4, and so on. Keep the ```json``` formatting intact.

Figure 15. Prompt template for descriptive QA pairs generation. We prompt GPT-4o to focus on various aspects of descriptive queries, such as textual content, numerical content, and graphical content, and generate multiple question-answer pairs based on these dimensions.

Prompt for Reasoning QA Pairs Generation

User:
As an expert in chart mathematical analysis, you are provided with the original chart code as follows:

```python
{chart_code}
```

You are tasked with examining a given chart code and generating a series of in-depth **Analytical Reasoning-based** question-answer pairs. Additionally, you have been given the chart image. The chart image will help you verify if the questions can be fully answered from it.

Generation Requirements:
1. Generate a diverse set of questions based on the chart data. The questions should cover a variety of aspects, such as comparisons, ratios, rankings, trends, range, overlaps, extreme values, relationships, highlights, annotations, shading area or calculations. You can freely combine any of above aspects to create a challenging question.

2. Ensure that the phrasing of the questions is varied and complex. Avoid using the same simple question structure or repetitive wording. Use a variety of question types, including but not limited to 'Is/Are', 'What', 'How', 'How many (times)', 'Which', 'Where', 'when' and 'If'. Incorporate different sentence structures, such as direct, indirect, conditional (e.g., 'If/Assuming ...', 'Excluding/Discarding ...', 'As ....'), rhetorical, **selective** (e.g., '... or ...?') and **negated** questions (e.g., 'Which is not ...?').

3. Include nested questions that require layered reasoning, as well as declarative sentences that prompt identification or clarification (e.g., 'Compared to/Comparing/Considering/Focusing on ...', 'In terms of/Given ...', 'Referring to ...', 'In/On/At/Across/Around/For/Of/Over/From/Between/By/Among/With/During/Towards ...', 'Based on ...', etc.). You can use a combination of one declarative instruction and one question, but must not to combine two questions.

4. Be sure to ask questions that address different dimensions of the data, such as time, categories, labels, names, values, and relative changes, rather than focusing only on the largest or smallest values (consider including **middle-ranked** values/labels such as 'the second', 'the third' etc., for a more balanced analysis).

5. If the chart contains subplots, prioritize questions that **involve reasoning across multiple subplots**, or focus on a specific subplot (by indicating its location, name, or ID) to explore detailed aspects of the data.

6. The generated questions should **exhibit clear visual references**, with an emphasis on the nuanced variations in details. **Be sure to ask how the legend labels or colorbar interact with the data.** The corresponding answers can be text or numbers, but not always number of values.

Response Format:
Please use the following JSON format to ensure clarity:

```json
{
  {
    "image_id": "{img_id}_1",
    "question": "Provide your question here",
    "rationale": "Provide your concise analysis and calculation procedures (MUST enumerate **all relevant elements** with **concrete values** and show **detailed calculations with specific values**)",
     "answer": "Provide complete final answer here",
     "rating": "Rate your confidence in this question, rationale and answer [0-5], Integer",
   },
   {
    "image_id": "{img_id}_2",
     ...
   },
     ...
}
```

Notes:
1. Focus only on aspects that can be inferred visually from the chart image. Do not refer to facts or terms that directly suggest the content is based on the code rather than the chart image itself.
2. Frame questions to encourage mathematical reasoning and logical deductions, avoiding sole reliance on chart data.
3. **MUST ENSURE** the rationale provides **numerical/text recognition or calculations**, converting visual questions into recognition- and calculation-based questions to support clear and logical reasoning.
4. The total number of questions should be **10 to 15**, and feel free to rephrase your questions to enhance diversity.
5. **MUST INCLUDE** one instance of "Not Applicable" for answer (different from "No") when appropriate. If a question relates to specific content that is missing in the image (especially for legend and colorbar, such as when asking about legend (e.g., no legend), values of colorbar (e.g., no colorbar)), answer "Not Applicable". Questions categorized as 'Not Applicable' should begin with 'What' and be generated based on your judgement.
6. Do not include questions about **alpha value, transparency, opacity, padding, spacing, grid, width, viewing angle, role, rotation, function (e.g.,'np.random'), font size, figsize, interpretation, visualization, significance, impact, save filename and save_path** in the code. Focus only on the meaningful visual elements in the chart and avoid mentioning 'code' in the questions, rationales and answers.
7. Ensure the image_id numbers increment sequentially, such as _1, _2, _3, _4, and so on. Keep the ```json``` formatting intact.

Figure 16. Prompt template for reasoning QA pairs generation. Prior to directly generating the final answer, we generate a rationale for in-depth analysis.

Prompt for Evaluating Model Predictions on ECDBench

Your task is to rigorously evaluate whether the VLM's prediction aligns with the expected answer for a question regarding a chart (note: the chart image itself is not provided here). The evaluation should focus solely on factual alignment between the prediction and the ground truth. Minor differences in wording or phrasing are acceptable as long as the core meaning remains consistent.

For numerical answers, **a margin of error up to 5% is acceptable** unless explicitly stated otherwise in the question. However, partial correctness or incomplete responses should not be considered correct.

- Question: {question}
- Expected Answer: {answer}
- Prediction: {prediction}

Please respond using the following format:
Correctness: (Yes or No)

Figure 17. Prompt template for evaluating model predictions on our constructed ECDBench. For numerical responses, we allow a tolerance range of 5% during evaluation.