

C²: Scalable Auto-Feedback for LLM-based Chart Generation

Woosung Koh^{1*}†, Jang Han Yoon^{1*}, MinHyung Lee¹, Youngjin Song¹,
Jaegwan Cho¹, Jaehyun Kang¹, Taehyeon Kim^{2‡}, Se-Young Yun^{2§},
Youngjae Yu^{1§}, Bongshin Lee^{1§}

¹Yonsei University, ²KAIST AI
{reiss.koh, jeffrobot99, yjy, b.lee}@yonsei.ac.kr
{potter32, yunseyoung}@kaist.ac.kr

Abstract

Generating high-quality charts with Large Language Models (LLMs) presents significant challenges due to limited data and the high cost of scaling through human curation. \langle instruction, data, code \rangle triplets are scarce and expensive to manually curate as their creation demands technical expertise. To address this scalability challenge, we introduce a *reference-free* automatic feedback generator, which eliminates the need for costly human intervention. Our novel framework, C², consists of (1) an automatic feedback provider (CHARTAF) and (2) a diverse, reference-free dataset (CHARTUIE-8K). The results are compelling: in our first experiment, 74% of respondents strongly preferred, and 10% preferred, the results after feedback. The second post-feedback experiment demonstrates that CHARTAF outperform nine baselines. Moreover, CHARTUIE-8K significantly improves data diversity by increasing queries, datasets, and chart types by 5982%, 1936%, and 91%, respectively, over benchmarks. Finally, a study of LLM users revealed that 94% of participants preferred CHARTUIE-8K’s queries, with 93% deeming them aligned with real-world use cases. Core contributions are available as open-source at chartsquared.github.io, with ample qualitative examples.

1 Introduction

Charts are a powerful means to convey information in diverse fields, including journalism, business, and scientific research (Fox and Hendler, 2011; Rodríguez et al., 2015; Islam and Jin, 2019a). With the success of foundation models (Kaplan et al., 2020; Roziere et al., 2023), there has been an increasing demand for generating charts using Large Language Models (LLMs). For instance, LIDA (Dibia,

2023) uses LLMs to automatically generate visualizations and infographics, and Chat2VIS (Maddigan and Susnjak, 2023) incorporates LLMs to create charts from natural language queries. Moreover, LLM-generated charts empower humans by helping non-experts generate high-quality charts (Maddigan and Susnjak, 2023) and improving accessibility to those with special needs (Gorniak et al., 2023; Moured et al., 2024). Despite the rising interest, two key challenges persist: (i) the difficulty in evaluating LLM-generated charts and (ii) the limited availability of training data.

(i) Chart generation lacks straightforward evaluation methods, making it difficult to assess and improve the quality of LLM-generated charts. Unlike tasks with clear-cut answers, such as mathematical problem-solving where verifiers can automatically assess correctness (Uesato et al., 2022; Wang et al., 2024a), chart evaluation is inherently subjective. Multiple correct designs may exist for a task (or goal), and quality often aligns with human aesthetic and functional preferences. Consequently, current evaluation systems (Yang et al., 2024; Wu et al., 2024; Xia et al., 2024) rely on reference-based approaches, necessitating labor-intensive \langle instruction, data, code¹ \rangle triplets as a gold reference for evaluation and thus limiting their scalability.

(ii) Furthermore, in contrast to image generation, which typically requires only \langle instruction, image \rangle pairs (Radford et al., 2021), chart generation demands more complex \langle instruction, data, code \rangle triplets. This significantly increases the costs associated with data collection and annotation. The limited number and diversity of available data restrict the variety of charts users can generate, making chart generation expensive and labor-intensive even for common applications (Vázquez, 2024).

*Equal contribution, co-first authors

†Work done while an intern at KAIST AI

‡Mentor

§Co-corresponding authors

¹Code here can be replaced with the image generated by executing the code.

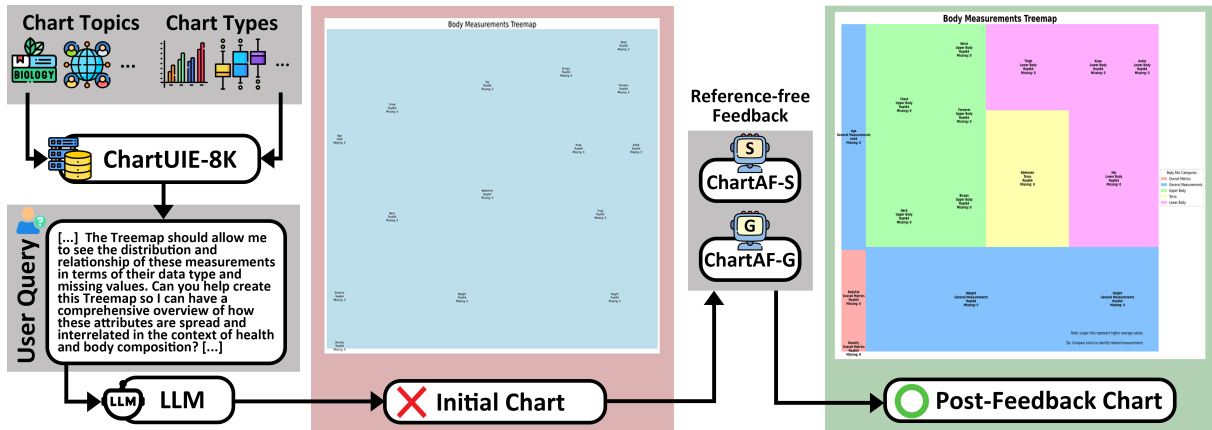


Figure 1: Schematic overview of C^2 illustrating the synergy between CHARTUIE-8K and CHARTAF. The scale is made possible by CHARTAF’s capability to provide *reference-free* feedback. An end-to-end example is available on our [project site](#)’s [github README](#).

To effectively tackle both (i) and (ii), it is essential to first address the primary bottleneck—reference-based evaluation—which also opens the door to significantly improving data diversity and scale. To this end, we introduce C^2 , a scalable framework, composed of the following two synergistic components:

- **CHARTAF**: A pipeline for Chart Auto-Feedback, comprising CHARTAF-S for evaluation scores and CHARTAF-G for granular feedback in natural language.
- **CHARTUIE-8K**: A large-scale (over 8,000 instances) Chart User Interaction Emulation dataset.

Fig. 1 illustrates a schematic overview of C^2 . CHARTAF empowers automatic (i.e., human-annotation-free) chart generation improvements (Sec. 2). CHARTAF works exceptionally well reference-free, enabling the cost-effective curation of a large-scale chart generation evaluation set CHARTUIE-8K (Sec. 3).

The quantitative and qualitative results of C^2 is overwhelmingly positive. First, by leveraging CHARTAF’s scalar evaluation scores for a simple test-time scaling scheme, 84% of respondents preferred the post-feedback results, with 74% strongly preferring them and 10% preferring them (Sec. 4.2). Second, employing CHARTAF’s granular feedback to in-context tune, CHARTAF’s post-feedback preference scores beat 9 baselines, that are alternatives to CHARTAF (Sec. 4.3). The qualitative improvements for both test-time scaling and in-context tuning can be viewed on our open-source [project site](#).

Finally, CHARTAF’s reference-free nature allows us to curate, CHARTUIE-8K, dramatically raising data diversity via number of queries, underlying datasets, and chart types by 5982%, 1936%, and 91%, respectively, against existing evaluation sets (Tab. 1). We also demonstrate that CHARTUIE-8K closely aligns with real-world human requests via a study of LLM users (Sec. 4.4). The study highlights that CHARTUIE-8K’s evaluation set distribution closely aligns with real-world users, and 94% and 93%, prefer, and think is realistic, respectively.

1.1 Related Work

Chart Generation. Maddigan and Susnjak (2023) offer prompt-engineered, LLM-based chart generation—however, they only provide qualitative case studies to verify their contribution. Dibia (2023) presents an LLM-based infographic visualization tool, which includes interactive charts. However, Dibia (2023) does not include a human study, making it challenging to assess its effectiveness. Sah et al. (2024) propose a natural language-to-chart recommendation approach based on the visualization language Vega-Lite. Therefore, their task deviates from the LLM-based chart generation we tackle. Tian et al. (2024) recently proposes a chart generation work that generates charts from "abstract user utterances" (as stated in the paper), which diverges from the instruction-based queries we address. An example of an "abstract user utterance" they provide is "What kind of movies earn the most recently?"—this diverges from the example instructions provided in our study of LLM users (App. B.2). This difference is understandable as

their chart generation is based on proprietary user-interaction software, not a common LLM chatbot.

Tuning with Feedback. Feedback-based LLM tuning is commonly done via three methods: parameter tuning (Ouyang et al., 2022), in-context tuning (Liu et al., 2022), and test-time scaling (TTS; OpenAI (2024)). Parameter tuning occurs through (implicit or explicit) rewards, requiring large amounts of reward data. For instance, Kirstain et al. (2023) and Xu et al. (2024b) collected over 500,000 and 137,000 annotations, respectively. In-context tuning allows LLMs to improve via n -shot generation (Chen et al., 2023), which requires n additional prompts. TTS refers to the case where a verifier with an ordinal output (typically of scalar value) can help improve the final generation (Snell et al., 2024). Each tuning method requires an external feedback-provider, such as a reward model, additional prompt, or verifier.

2 C²: CHARTAF

In C², a feedback-provider, CHARTAF enables TTS and in-context tuning. We first describe the shortcomings of existing approaches, and introduce two versions of CHARTAF: CHARTAF-S and CHARTAF-G.

2.1 Towards High-performing Feedback

Two feedback types have been explored in the LLM research community: (1) scalar score-based ($s_N \in \mathbb{N}_{\cup 0}$) and (2) natural language-based ($s_N \in \mathcal{N}$, where \mathcal{N} is the natural language set space) feedback. Instead of competing, (1) and (2) serve different purposes: (1) is suited for parameter tuning with rewards, and TTS, while (2) is suited for in-context tuning. We summarize prior efforts of the two approaches in chart generation below.

(1) Yang et al. (2024), Wu et al. (2024), and Xia et al. (2024) provide $[0, 100]$, $[1, 10]$, $[0, 5]$, scores, respectively. However, as these methods were developed under a reference-based regime, their performance sharply deprecates when applied reference-free. While there are other similar works, they are closed-source without clear details for replication (Han et al., 2023; Xu et al., 2023).

(2) A naïve approach uses a zero-shot LLM to provide feedback (Yang et al., 2024), which we call Naïve Feedback (NF). This method, as our experiments later show, is ineffective (Tab. 3). Furthermore, Yang et al. (2024) does not provide human

studies of their NF against baselines so there is no evidence on its efficacy.

To overcome the limitations of past works, we present CHARTAF (Fig. 2). As CHARTAF-S is a subset of CHARTAF-G, we first introduce CHARTAF-S and then the additional component corresponding to CHARTAF-G. Fine-grained pseudocode and prompts are provided in App. A.

2.2 CHARTAF-S (\tilde{f}_{AF})

Module 1. The user query, $q \in \mathcal{Q} \in \mathcal{N}$, is first decomposed into three essential factors in a chart generation query: Task, Purpose, and Audience (Choe et al., 2017; Lee et al., 2020; Narechania et al., 2020; Wang et al., 2020; Parsons, 2021). By explicitly decomposing the q into Task, Purpose, and Audience, CHARTAF can better infer the user intention (Quadri et al., 2024; Bressa et al., 2024). Since user queries are often brief (Fig. 6, App. B.2), the intention may not always be clearly stated. Nevertheless, CHARTAF utilizes the q to induce underlying intentions.

This information is then fused with the Basic Criteria—a general, high-level criteria applicable to all chart evaluations. The Basic Criteria ensures that CHARTAF comprehensively considers chart elements: Chart Type, Visual Embellishment, Text, Color, Annotation, Aesthetics, and Visual Clutter. This criteria is inspired by the rich literature in visualization research. We document the research that corresponds to each element in App. A.1.

This first module is the domain grounding module. CHARTAF replaces gathering costly human annotations with existing scholarly research. Not only is this cheaper, this approach closely aligns with how human-made chart generations would be evaluated. We would grade human students’ chart generations with a domain expert (lecturer) that has learned the principles of chart generation, grounded in scholarly literature (Bach et al., 2023).

Module 2. Considering the decomposed q , and Basic Criteria, CHARTAF generates q -Specific Criteria—specializing the general Basic Criteria. Specialization is key to generating feedback that is customized to the specific q (Kim et al., 2024). The criteria is then transformed to binary (i.e., yes or no) questions (Hu et al., 2023) as we find that LLMs are more reliable when reasoning with binary rather than open-ended questions. This also allows CHARTAF to explicitly associate each criterion with one question. Otherwise, LLMs tend

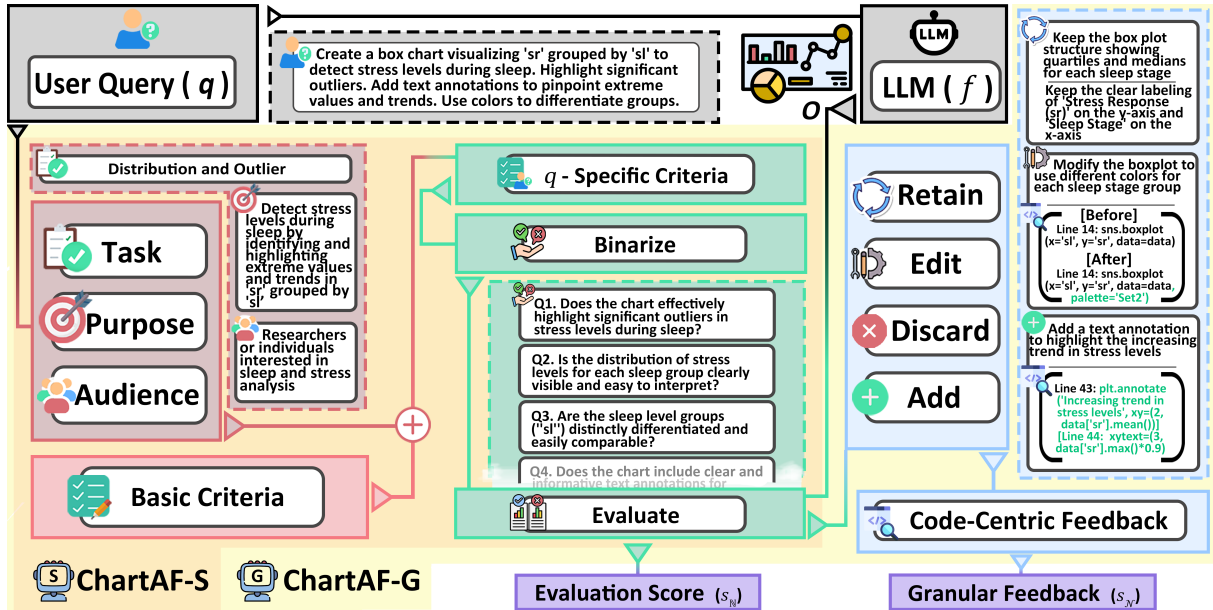


Figure 2: Schematic diagram of CHARTAF, including a qualitative example indicated by dashed containers. The process starts in the top-left with the user query, which is processed by the chart-generating LLM and either CHARTAF-S or CHARTAF-G. Notably, CHARTAF-S and CHARTAF-G both share the first two modules (red and green). The final output is a scalar evaluation score or granular feedback, depending on the chosen path.

to lump numerous criterion together, resulting in duplicate related criteria.

Then, these questions are evaluated considering the LLM-generated chart (generated by executing the code output $o \in \mathcal{O} \in \mathcal{N}$). This code is generated by the chart generating LLM, $f : \mathcal{Q} \mapsto \mathcal{O}$. The final scalar, s_N , is derived by equal-weighting each binary answer, yes (1) and no (0). If researchers only require a scalar score, for downstream applications, the process can be terminated.

2.3 CHARTAF-G (f_{AF})

Module 3. For granular feedback, the process continues by associating each answered question with Retain, Edit, Discard, and Add. This association helps to decompose the evaluation result of each criterion to *actionable* feedback. Based on this association, code-centric text feedback is provided (Bi et al., 2024): the LLM is prompted to provide fine-grained feedback considering the downstream application, i.e., chart generation via code generation. Code-centric feedback encourages the LLM to provide explicit feedback that can be directly applied to the downstream f .

3 C²: CHARTUIE-8K

Leveraging the reference-free nature of CHARTAF, we curate CHARTUIE-8K, a comprehensive chart generation evaluation set. Since no gold references

are required, we can significantly scale the dataset. Qualitative examples of CHARTUIE-8K can be viewed in Figs. 1, 2, App. C, and our project site.

Curation Method. An overview of CHARTUIE-8K’s curation process is illustrated in Fig. 3. First, for diverse chart topics, we semi-automatically crawl diverse datasets online with appropriate licenses. We include only the datasets that have been used by humans for chart generation purposes, as not all datasets are suitable for visualization. Next, to ensure diverse chart types and annotations, we adopt a comprehensive list of chart types (Hess, 2022) and annotations (Ren et al., 2017). Then, we emulate two types of users: [U1] lay users and [U2] detailed users.

To this end, we synthetically (LLM-assisted) generate initial instructions with two configurations: approx. [U1] <50, and [U2] <100 words. Considering the salience of multi-turn benchmarking (Wang et al., 2024c), we further emulate a single QA cycle. I.e., the LLM asks for clarifying questions to the user, then, the user emulator responds to [U1] 25% or [U2] 50% of the questions. The pseudocode is provided in App. B.1.

Statistical Summary. To systematically understand the evaluation set, we compare key statistics against relevant benchmarks (see Tab. 1). We exclude chart topic counts for benchmarks, as they are

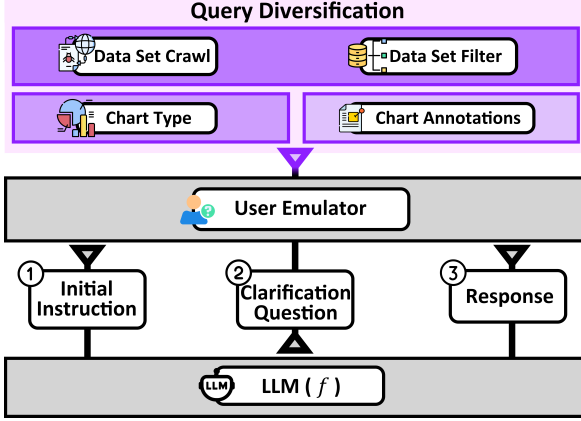


Figure 3: CHARTUIE-8K curation schematic diagram.

Number of	CHARTUIE-8K	MatPlotBench	Plot2Code
Queries	8028 (+5982%)	100	132
Datasets	509 (+1936%)	25	2
Chart topics	44	—	—
Chart types	63 (+91%)	33	24 (6)

Table 1: Comparing key statistics of chart generation evaluation sets. **Bold** represents improvement from the best existing benchmark.

not provided in MatPlotBench (Yang et al., 2024) and Plot2Code (Wu et al., 2024). We do not manually count their chart topics as this a subjective task. On the other hand, we count MatPlotBench’s and Plot2Code’s chart types using the same taxonomy used for CHARTUIE-8K (App. B.1). We leave the original number of coarse chart types provided by Plot2Code in parentheses for documentation. Finally, the distribution of chart topics and types of CHARTUIE-8K are presented in Fig. 4.

4 Empirical Study

4.1 Preliminary and Notations

Denote the chart code-generating LLM, $f : \mathcal{Q} \mapsto \mathcal{O}$. Let a feedback-provider take the output of f , including the executed chart image, then $\tilde{f}_{AF} : \mathcal{O} \mapsto \mathcal{S}_N \in \mathbb{N}_{\cup 0}$, $f_{AF} : \mathcal{O} \mapsto \mathcal{S}_N \in \mathcal{N}$, for TTS, and in-context tuning feedback, respectively. Following Liang et al. (2024), $h : \mathcal{O}_{pre} \times \mathcal{O}_{post} \mapsto \{-2, -1, 0, 1, 2\}$ represents the human post-feedback preference score function. The co-domain refers to \succ : strongly prefer pre-feedback (-2), \succeq : prefer pre-feedback (-1), \sim : indifferent (0), \preceq : prefer post-feedback (1), \prec : strongly prefer post-feedback (2). To demonstrate CHARTAF-S (\tilde{f}_{AF}) and CHARTAF-G (f_{AF})’s utility, we empirically study TTS with CHARTAF-S (Sec. 4.2) and in-context tuning with CHARTAF-G (Sec. 4.3).

\succ	\succeq	\sim	\preceq	\prec
3 (3.9%)	4 (5.2%)	5 (6.5%)	8 (10.4%)	57 (74%)

Table 2: Number of respondents out of 77 (%). \succ : strongly prefer pre-TTS, \succeq : prefer pre-TTS, \sim : indifferent, \preceq : prefer post-TTS, \prec : strongly prefer post-TTS.

Fine-grained details are reported in App. D.

4.2 Test-time Scaling with CHARTAF

Experiment Set-up. We demonstrate that CHARTAF-S is an effective verifier for TTS. Following Snell et al. (2024)’s parallel best-of- N , we generate $N \in \mathbb{N}$ independent samples and choose the one with the highest score ($s_N \in \mathcal{S}_N := [0, 100]$). Here, the unit of inference budget is Ns . Let memory set $\mathcal{M} := \{o_1, \dots, o_N\}$, $o \in \mathcal{O}$, hold the N independent outputs of f . Then, we experimentally show that

$$\hat{o} := \arg \max_{o \in \mathcal{M}} \tilde{f}_{AF}(o), \quad (1)$$

$$\begin{aligned} & \text{if } \hat{o}_i > \hat{o}_j, i \neq j \in \mathbb{N}, \\ & \Rightarrow h(\hat{o}_{j:=1} := f(q), \hat{o}_i) > 0, \end{aligned} \quad (2)$$


across $q \in \mathcal{Q}$.

We employ a double-blind 120 human study (of which 77 pass the rigorous sanity check) to compare pre-TTS ($N := 1$) and post-TTS ($N := 4$) preference scores. Each participant is presented with a random TTS sample. For this experiment we set both f and \tilde{f}_{AF} backbone as GPT 4o to emphasize that \tilde{f}_{AF} does not have to be a superior model for CHARTAF-S to be useful.


Results. First, the scaling curve is depicted in Fig. 5. The positive slope highlights that CHARTAF can be effectively used as a TTS verifier. To rigorously verify this claim, the distribution of the human study is presented in Tab. 2. Going from a median s_N of 40.47→62.75 (pre→post-TTS) leads to 74% strongly preferring the post-TTS and 10.4% preferring post-TTS. This is a strong indication that s_N is closely proxying human preferences.

4.3 In-context Tuning with CHARTAF

Experiment Set-up. For a comprehensive empirical study we consider four LLMs: (i) GPT 4o (Achiam et al., 2023), (ii) Claude 3.5 Sonnet (Anthropic, 2024), (iii) Llama 3.1 70b (Dubey et al., 2024), and (iv) Gemma 2 27B (Team et al.,



(a) Chart topics in order of frequency



(b) Chart types in order of frequency

Figure 4: CHARTUIE-8K distribution. Top 10 topics and types are explicitly depicted while the remaining is classified as others. (n) is the number of samples out of 8028.




Figure 5: Test-time scaling with CHARTAF as a verifier. Raising N leads to improved generations.

2024)—of which (i) and (ii) are closed-source and (iii) and (iv) are open-source. The two closed-source models are used as f_{AF} backbones, while all four models are used as f . Detailed LLM configurations are provided in App. E.

Our empirical study employs a double-blind human study of 60 queries with 120 participants, 77 of whom pass the rigorous sanity check. 120 individuals are required for 60 queries as we test the two f_{AF} independently. Within the 60 queries, 15 are designated for each of the four f . For representative sampling, we random sample the 15 queries with two constraints: (1) each of the 15 queries ask for different chart types, (2) 1:1 ratio of [U1] and [U2] queries. Concretely, our experiments show

$$\begin{aligned}
 &h(f(q), f(q \oplus f_{AF}(q))) > \\
 &h(f(q), f(q \oplus f_{AF}^{baseline}(q))) \quad (3)
 \end{aligned}$$

across $q \in \mathcal{Q}$, where \oplus denotes string concatenation.

Baselines. We identify four baselines: (1) ChartX (Xia et al., 2024), (2) MatPlotBench (Yang et al.,

2024), (3) Plot2Code (Wu et al., 2024), and (4) ChatEval (Chan et al., 2024) that can provide feedback for chart generation. (1) (2) (3) are chart generation specific feedback providers, while (4) is a generalist. One advantage of (1) (2) (3) is its token cost light nature. Therefore, for a fair comparison vis-à-vis token cost, we enhance these baselines with Auto-Chain-of-Thought (A-CoT) (Zhang et al., 2023) and Self-Consistency (SC) (Wang et al., 2023). Beforehand, we ran preliminary studies with A-CoT, SC, and Self-Refine (Madaan et al., 2023), and found that A-CoT+SC performed best. To solidify the utility of CHARTAF we include two additional baselines. The first is Naïve Feedback (NF) which is zero-shot asking another LLM to give feedback (Yang et al., 2024). The second is skipping the feedback stage, and instead directly adding A-CoT to f (Skip+A-CoT).

Results. We present our findings in Tab. 3 including (input plus output) token costs. CHARTAF ranks **first** 13 (out of 16) times and **second** for the remaining three. This level of consistency across four different models accentuates CHARTAF’s utility across large and smaller f . Furthermore, regardless of f_{AF} backbone, post-feedback always (8 out of 8) results in improvement, $\mu > 0$, spotlighting CHARTAF’s universality. Detailed qualitative examples are provided on the [project site](#).

It is important to note that CHARTAF’s distinguished performance cannot be trivially matched by enhancing baselines with state-of-the-art prompting methods. While the three enhanced baselines use **128%** (ChartX+A-CoT+SC), **85%** (MatPlotBench+A-CoT+SC), **70%** (Plot2Code+A-CoT+SC) more tokens on average than CHARTAF, they fail to be competitive. At the time of ex-

Table 3: Empirical study preference scores as presented in Sec. 4.1. μ represents mean, $\tilde{\mu}$ represents median. **Bold** is the best result across the column, and underlined is the second-best. Total (input plus output) token costs per query evaluation are reported as $\mu \pm \sigma$.

f	GPT 4o		Claude 3.5 Sonnet		Llama 3.1 70b		Gemma 2 27B		f_{AF}
Statistic	μ	$\tilde{\mu}$	μ	$\tilde{\mu}$	μ	$\tilde{\mu}$	μ	$\tilde{\mu}$	Token Cost
f_{AF} Backbone	GPT 4o								
CHARTAF (ours)	0.5	1	0.5	1.5	1.308	2	0.333	1	14013 \pm 1155
ChatEval	-0.375	-1	-0.583	-0.5	-0.615	-1	-1.222	-2	70195 \pm 8121
ChartX+A-CoT+SC	-0.25	-1	-0.167	<u>0</u>	-1.154	-2	<u>0.111</u>	1	31791 \pm 4604
MatPlotBench+A-CoT+SC	-0.25	<u>0</u>	-0.5	-1	-0.231	<u>0</u>	-0.556	-1	24773 \pm 2149
Plot2Code+A-CoT+SC	-0.75	-1.5	-0.25	<u>0</u>	-0.385	<u>0</u>	-1	-1	22519 \pm 2122
ChartX	<u>0.5</u>	<u>0.5</u>	<u>0.083</u>	<u>1</u>	-0.385	<u>0</u>	-1.111	-2	1665 \pm 235
MatPlotBench	0.25	<u>0.5</u>	-0.5	-1	-0.385	<u>0</u>	-0.444	<u>0</u>	1653 \pm 284
Plot2Code	-0.5	-0.5	-0.167	-0.5	-0.077	<u>0</u>	-0.667	-1	1604 \pm 277
NF	0.75	1	-0.583	-1	-0.692	-2	-0.333	-2	2739 \pm 385
Skip+A-CoT	-1.375	-2	-1	-1	-1.692	-2	-1.667	-2	—
f_{AF} Backbone	Claude 3.5 Sonnet								
CHARTAF (ours)	0.167	<u>0</u>	<u>0.273</u>	1	1.273	2	1	2	14927 \pm 981
ChatEval	<u>0</u>	0.5	-1	-1	<u>0</u>	<u>0</u>	<u>0.143</u>	<u>1</u>	74369 \pm 8397
ChartX+A-CoT+SC	-0.667	-1	-0.727	-2	<u>0.818</u>	<u>1</u>	-0.571	-1	34210 \pm 4084
MatPlotBench+A-CoT+SC	-1.333	-1.5	<u>0.273</u>	1	-0.455	-1	-1	-2	28980 \pm 2608
Plot2Code+A-CoT+SC	-0.667	-1.5	-0.727	-1	<u>0.273</u>	<u>1</u>	-0.286	<u>0</u>	26849 \pm 2520
ChartX	-0.167	<u>0</u>	-0.364	<u>0</u>	-0.545	-1	-1.143	-1	1753 \pm 277
MatPlotBench	-0.5	-0.5	0.364	<u>0</u>	-0.545	-1	<u>0</u>	<u>0</u>	2040 \pm 407
Plot2Code	-1.167	-1.5	-0.273	<u>0</u>	<u>0.182</u>	<u>0</u>	-0.143	<u>0</u>	1851 \pm 412
NF	-1.167	-1.5	-1.727	-2	<u>0</u>	<u>1</u>	-1.571	-2	2739 \pm 562
Skip+A-CoT	-1	-1	-1.182	-2	-1.636	-2	-1.571	-2	—

Color-coded: $\mu, \tilde{\mu} \leq -1.5$; $-1.5 < \mu, \tilde{\mu} < 0$; $\mu, \tilde{\mu} = 0$; $0 < \mu, \tilde{\mu} < 1.5$; $\mu, \tilde{\mu} \geq 1.5$

perimentation, an evaluation for a single query costed \$0.1 and \$0.12 on average using GPT 4o and Claude 3.5 Sonnet, respectively.

4.4 CHARTUIE-8K Experiments

Experiment Set-up. We provide evidence than our novel evaluation set is of high-quality, and potentially more realistic than past benchmarks. We employ a double-blind human study of 130 participants (89 of whom pass the rigorous sanity check). Each participant is given a randomized chart image they are imagining, and asked to prompt an LLM their initial instructions. As the factors that comprise query realism are qualitative in nature, we analyze three dimensions that can be quantitatively captured: (i) the word count distribution, (ii) % of respondents preferring CHARTUIE-8K’s interaction, and (iii) % of respondents who think CHARTUIE-8K’s user emulation is realistic. Lastly, the participant is asked whether the extra QA cycle is desirable and realistic when interacting with an LLM. Details of the human study is provided in App. D.4.

Results. We visualize the results in Fig. 6. In terms of word count, CHARTUIE-8K (**green**) most closely matches the ground truth distribution of the study of LLM users (**purple**). Please see some qualitative examples of the instructions the users gave in App. B.2. Notably, 94% of respondents prefer the extra QA cycle, and 93% believe CHARTUIE’s user emulation is realistic. We also qualitatively observe, as presented in App. C, existing evaluation sets’ queries are too technical to be realistic.

5 Discussion and Impact of C²

5.1 CHARTAF Enables Scalable Feedback

Scalability. As described in Sec. 1, the lack of training data, $\langle \text{instruction, data, code} \rangle$, can be mitigated by an effective reference-free feedback provider that only requires $\langle \text{instruction, data} \rangle$. However, as indicated by the predominance of **red** in Tab. 3, existing feedback providers perform poorly under this reference-free regime. Furthermore, increasing token usage via current methods fails to resolve this issue. CHARTAF addresses this scalability barrier, as demonstrated by the experimental results in in-context tuning and TTS.




Figure 6: Results of the CHARTUIE-8K empirical study. **Top:** Ground truth Q^* (purple) vs. Q distributions of initial instruction word count. **Middle:** % of respondents preferring CHARTUIE’s interaction. **Bottom:** % of respondents who think CHARTUIE’s user emulation is realistic.

In-context Tuning with CHARTAF. CHARTAF performs best among existing methods for nearly all f and f_{AF} (Tab. 3). This is particularly notable, as in-context tuning is often challenging for smaller models. Even when the feedback is helpful, the intrinsic limitations of f —such as small model size or insufficient tuning to follow instructions—may result in f inaccurately reflecting the feedback. Despite these challenges, the granular feedback of CHARTAF effectively improves smaller models (Llama 3.1 70b, Gemma 2 27b). In fact, Llama 3.1 70b and Gemma 2 27b on average experience greatest performance improvements after in-context tuning with CHARTAF, likely due to their weaker base performance that allows for larger post-tuning gains.

Towards Complex Task Verification. Additionally, CHARTAF’s reference-free TTS performance shows significant strength. TTS has been a paradigm-shifting development as it introduces a novel neural scaling axis. While previous scaling laws focused on data and training compute, recent works show that similar scaling laws apply to the inference compute axis. However, current TTS is limited to tasks with reliable and cheap verifiers (Wang et al., 2024a,b), emphasizing the salience of fast, reliable, and cost-effective verifiers (Brown et al., 2024; Snell et al., 2024). CHARTAF is the first effective demonstration of a chart generation

verifier within the TTS framework.

We encourage future works to build upon our approach to advance TTS for chart generation and other similar tasks. We used the simplest approach for TTS to empirically prove the verifier’s effectiveness. More advanced TTS for chart generation is an open problem.

Scaling Frontier Models with CHARTAF. Notably, CHARTAF is not a distillation method transferring knowledge from larger→smaller models. In Tab. 3, CHARTAF remains effective even for $\langle f, f_{AF} \rangle$ pairs where f and f_{AF} are similarly performing models. Furthermore, TTS is demonstrated using the same backbone LLM. Such self-improving LLMs are indispensable in advancing frontier models (Huang et al., 2023).

Scaling without Parameter Updates. We demonstrate the effectiveness of C^2 with no parameter updates. This is notable as parameter updating LLMs incur a large memory and training throughput cost. Therefore, C^2 is orthogonal to Zadeh et al. (2024), where they focus on automating the instruction (parameter) tuning process for chart generations.

5.2 Unlocking Large-scale Data with C^2

Comprehensive Evaluation and Generation. The strong reference-free performance of CHARTAF (Sec. 5.1) enables the curation of a cost-effective and diverse query set, CHARTUIE-8K. This approach contrasts with the reference-based nature of existing query sets, which suffer from limited diversity (Tab. 1). By pairing CHARTAF with CHARTUIE-8K, the combined framework of C^2 supports a *broadly* inclusive evaluation set (Tab. 1, Fig. 4). This enables researchers to evaluate models on a more comprehensive and diverse set of queries. Moreover, through CHARTAF, C^2 can generate large-scale, high-quality outputs that significantly improve over previous methods.

Realistic Evaluation. Lastly, it is crucial to create evaluation sets that closely reflect real-world use-cases. As shown by the distributions in Fig. 6, existing query sets often diverge significantly from common user queries, reducing their practical utility. We recommend that future work proposing evaluation sets include rigorous studies (e.g. Sec. 4.4) to ensure their assets are pragmatically aligned with real-world use-cases.

6 Limitations

We discuss the limitations of this work to ensure full transparency. To our knowledge, we disclose all reasonable information throughout the paper, project site, and github repository.

Code Execution Error. The feedbacks, s_N , $s_{\mathcal{N}}$, presented in this paper is conditioned on the fact that the chart image has been successfully generated from the code. Occasionally, the code fails to execute or executes with an error. Under such circumstances, we allow a maximum of 5 re-generations for the initial f inference, and a maximum of 3 re-generations for post-feedback f inference. We document the initial f inference error rate for each f in App. F.

Coverage. The coverage of CHARTUIE-8K and CHARTAF is limited to the English language. Additionally, this study is not conducted with smaller models, e.g. 8B parameter size LLMs. We leave investigating expanded coverage to future works.

7 Human Study Ethical Consideration

To our knowledge, we follow best practices in computer science human studies (Müller et al., 2014; Müller and Sedley, 2015). First, to guarantee the privacy of both researchers and surveyors the study is double-blind, and do not collect any unnecessary data. Second, we clearly indicate at the very start that this is an academic survey for an academic paper. We do not upload the raw data on the public domain to avoid any potential unethical usage. Third, surveyors voluntarily conduct the surveys and can choose to leave at any time. Fourth, we ensure that the surveyors are compensated fairly. While we do not include any geographic restrictions, we pay \$14 per hour, well above the U.S. federal minimum wage of \$7.25 (Henderson, 2024) as of the surveys. We pay surveyors within 48 hours of completing the survey. Finally, our surveys do not contain explicit or triggering content.

Acknowledgement

We would like to thank Yun Ho Ro for his assistance with the figures and Jian Kim for her help in preparing the human study survey. This work was supported in part by the Yonsei University Research Fund of 2024-22-0058, the National Research Foundation of Korea (NRF) Grant funded by the Korea government (MSIT) (RS-2024-00353125), and the Institute of Information

and Communications Technology Planning and Evaluation (IITP) Grant, Artificial Intelligence Graduate School Program, Yonsei University (RS-2020-II201361) and KAIST (RS-2019-II190075).

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Kiran Ajani, Elsie Lee, Cindy Xiong, Cole Nussbaumer Knaflig, William Kemper, and Steven Franconeri. 2022. [Declutter and focus: Empirically evaluating design guidelines for effective data communication](#). *IEEE Transactions on Visualization and Computer Graphics*, 28(10):3351–3364.
- Tiffany Andry, Christophe Hurter, François Lambotte, Pierre Fastrez, and Alexandru Telea. 2021. Interpreting the effect of embellishment on chart visualizations. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–15.
- Anthropic. 2024. The claude 3 model family: Opus, sonnet, haiku.
- Benjamin Bach, Mandy Keck, Fateme Rajabiyazdi, Tatiana Losev, Isabel Meirelles, Jason Dykes, Robert S Laramee, Mashael AlKadi, Christina Stoiber, Samuel Huron, et al. 2023. Challenges and opportunities in data visualization education: A call to action. *IEEE Transactions on visualization and computer graphics*.
- Scott Bateman, Regan L Mandryk, Carl Gutwin, Aaron Genest, David McDine, and Christopher Brooks. 2010. Useful junk? the effects of visual embellishment on comprehension and memorability of charts. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 2573–2582.
- Zhangqian Bi, Yao Wan, Zheng Wang, Hongyu Zhang, Batu Guan, Fangxin Lu, Zili Zhang, Yulei Sui, Hai Jin, and Xuanhua Shi. 2024. [Iterative refinement of project-level code context for precise code generation with compiler feedback](#). *Preprint*, arXiv:2403.16792.
- Rita Borgo, Alfie Abdul-Rahman, Farhan Mohamed, Philip W Grant, Irene Reppa, Luciano Floridi, and Min Chen. 2012. An empirical study on using visual embellishments in visualization. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2759–2768.
- Nathalie Bressa, Jordan Louis, Wesley Willett, and Samuel Huron. 2024. Input visualization: Collecting and modifying data with visual representations. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 1–18.
- Bradley Brown, Jordan Juravsky, Ryan Ehrlich, Ronald Clark, Quoc V Le, Christopher Ré, and Azalia Mirhoseini. 2024. Large language monkeys: Scaling inference compute with repeated sampling. *arXiv preprint arXiv:2407.21787*.
- Nick Cawthon and Andrew Vande Moere. 2007. The effect of aesthetic on the usability of data visualization. In *2007 11th International Conference Information Visualization (IV'07)*, pages 637–648. IEEE.
- Chi-Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, and Zhiyuan Liu. 2024. [Chateval: Towards better LLM-based evaluators through multi-agent debate](#). In *The Twelfth International Conference on Learning Representations*.
- Siming Chen, Jie Li, Gennady Andrienko, Natalia Andrienko, Yun Wang, Phong H Nguyen, and Cagatay Turkyay. 2018. Supporting story synthesis: Bridging the gap between visual analytics and storytelling. *IEEE transactions on visualization and computer graphics*, 26(7):2499–2516.
- Yongrui Chen, Shenyu Zhang, Guilin Qi, and Xin-nan Guo. 2023. [Parameterizing context: Unleashing the power of parameter-efficient fine-tuning and in-context tuning for continual table semantic parsing](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 17795–17810. Curran Associates, Inc.
- Ming-Te Chi, Shih-Syun Lin, Shiang-Yi Chen, Chao-Hung Lin, and Tong-Yee Lee. 2015. Morphable word clouds for time-varying text data visualization. *IEEE transactions on visualization and computer graphics*, 21(12):1415–1426.
- Eun Kyoung Choe, Bongshin Lee, Haining Zhu, Nathalie Henry Riche, and Dominikus Baur. 2017. [Understanding self-reflection: how people reflect on personal data through visual data exploration](#). In *Proceedings of the 11th EAI International Conference on Pervasive Computing Technologies for Healthcare, PervasiveHealth '17*, page 173–182, New York, NY, USA. Association for Computing Machinery.
- Victor Dibia. 2023. Lida: A tool for automatic generation of grammar-agnostic visualizations and infographics using large language models. In *The 61st Annual Meeting Of The Association For Computational Linguistics*.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Geoffrey Ellis and Alan Dix. 2007. A taxonomy of clutter reduction for information visualisation. *IEEE transactions on visualization and computer graphics*, 13(6):1216–1223.
- Ana Figueiras. 2013. A typology for data visualization on the web. In *2013 17th International conference on information visualisation*, pages 351–358. IEEE.
- Peter Fox and James Hendler. 2011. Changing the equation on scientific data visualization. *Science*, 331(6018):705–708.

- Joshua Gorniak, Jacob Ottiger, Donglai Wei, and Nam Wook Kim. 2023. Vizability: Multimodal accessible data visualization with keyboard navigation and conversational interaction. In *Adjunct Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, pages 1–3.
- Arnav Gudibande, Eric Wallace, Charlie Victor Snell, Xinyang Geng, Hao Liu, Pieter Abbeel, Sergey Levine, and Dawn Song. 2024. [The false promise of imitating proprietary language models](#). In *The Twelfth International Conference on Learning Representations*.
- Yucheng Han, Chi Zhang, Xin Chen, Xu Yang, Zhibin Wang, Gang Yu, Bin Fu, and Hanwang Zhang. 2023. Chartllama: A multimodal llm for chart understanding and generation. *arXiv preprint arXiv:2311.16483*.
- Lane Harrison, Katharina Reinecke, and Remco Chang. 2015. Infographic aesthetics: Designing for the first impression. In *Proceedings of the 33rd Annual ACM conference on human factors in computing systems*, pages 1187–1190.
- Christopher G Healey. 1996. Choosing effective colours for data visualization. In *Proceedings of Seventh Annual IEEE Visualization'96*, pages 263–270. IEEE.
- Kaitlyn Henderson. 2024. The crisis of low wages: Who earns less than \$17 an hour in the us in 2024?
- Kosma Hess. 2022. 80 types of charts & graphs for data visualization (with examples). In *datylon*.
- Yushi Hu, Benlin Liu, Jungo Kasai, Yizhong Wang, Mari Ostendorf, Ranjay Krishna, and Noah A Smith. 2023. Tifa: Accurate and interpretable text-to-image faithfulness evaluation with question answering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20406–20417.
- Jiixin Huang, Shixiang Gu, Le Hou, Yuexin Wu, Xuezhi Wang, Hongkun Yu, and Jiawei Han. 2023. [Large language models can self-improve](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1051–1068, Singapore. Association for Computational Linguistics.
- Mohaiminul Islam and Shangzhu Jin. 2019a. [An overview of data visualization](#). In *2019 International Conference on Information Science and Communications Technologies (ICISCT)*, pages 1–7.
- Mohaiminul Islam and Shangzhu Jin. 2019b. An overview of data visualization. In *2019 International Conference on Information Science and Communications Technologies (ICISCT)*, pages 1–7. IEEE.
- Daekyoung Jung, Wonjae Kim, Hyunjoo Song, Jeong-in Hwang, Bongshin Lee, Bohyoung Kim, and Jinwook Seo. 2017. ChartSense: Interactive data extraction from chart images. In *Proceedings of the 2017 chi conference on human factors in computing systems*, pages 6706–6717.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.
- Dae Hyun Kim, Vidya Setlur, and Maneesh Agrawala. 2021. Towards understanding how readers integrate charts and captions: A case study with line charts. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–11.
- Hyoyoung Kim and Jin Wan Park. 2013. Topics on aesthetic data visualization: viewpoints, interpretation, and alternative senses. In *SIGGRAPH Asia 2013 Art Gallery*, pages 1–7.
- Tae Soo Kim, Yoonjoo Lee, Jamin Shin, Young-Ho Kim, and Juho Kim. 2024. [Evallm: Interactive evaluation of large language model prompts on user-defined criteria](#). In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI '24, New York, NY, USA. Association for Computing Machinery.
- Yuval Kirstain, Adam Polyak, Uriel Singer, Shahbuland Matiana, Joe Penna, and Omer Levy. 2023. Pick-a-pic: An open dataset of user preferences for text-to-image generation. *Advances in Neural Information Processing Systems*, 36:36652–36663.
- Primož Lavrič, Ciril Bohak, and Matija Marolt. 2017. Collaborative view-aligned annotations in web-based 3d medical data visualization. In *2017 40th International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*, pages 259–263. IEEE.
- Bongshin Lee, Eun Kyoung Choe, Petra Isenberg, Kim Marriott, and John Stasko. 2020. Reaching broader audiences with data visualization. *IEEE computer graphics and applications*, 40(2):82–90.
- Bongshin Lee, Nathalie Henry Riche, Petra Isenberg, and Sheelagh Carpendale. 2015. More than telling a story: Transforming data into visually shared stories. *IEEE computer graphics and applications*, 35(5):84–90.
- Sungkil Lee, Mike Sips, and Hans-Peter Seidel. 2012. Perceptually driven visibility optimization for categorical data visualization. *IEEE Transactions on visualization and computer graphics*, 19(10):1746–1757.
- Youwei Liang, Junfeng He, Gang Li, Peizhao Li, Arseniy Klimovskiy, Nicholas Carolan, Jiao Sun, Jordi Pont-Tuset, Sarah Young, Feng Yang, et al. 2024. Rich human feedback for text-to-image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19401–19411.
- Xiao Liu, Kaixuan Ji, Yicheng Fu, Weng Tam, Zhengxiao Du, Zhilin Yang, and Jie Tang. 2022. P-tuning: Prompt tuning can be comparable to fine-tuning

- across scales and tasks. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 61–68.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhunoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. 2023. [Self-refine: Iterative refinement with self-feedback](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Paula Maddigan and Teo Susnjak. 2023. [Chat2VIS: Generating data visualizations via natural language using chatgpt, codex and gpt-3 large language models](#). *IEEE Access*, 11:45181–45193.
- Stephen R Midway. 2020. Principles of effective data visualization. *Patterns*, 1(9).
- Omar Moured, Sara Alzababny, Anas Osman, Thorsten Schwarz, Karin Müller, and Rainer Stiefelhagen. 2024. Chartformer: A large vision language model for converting chart images into tactile accessible svgs. In *International Conference on Computers Helping People with Special Needs*, pages 299–305. Springer.
- Hendrik Müller, Aaron Sedley, and Elizabeth Ferrall-Nunge. 2014. Survey research in hci. *Ways of Knowing in HCI*, pages 229–266.
- Hendrik Müller and Aaron Sedley. 2015. [Designing surveys for hci research](#). In *CHI '15 Extended Abstracts on Human Factors in Computing Systems*, pages 2485–2486, New York, NY, USA.
- Arpit Narechania, Arjun Srinivasan, and John Stasko. 2020. NL4DV: A toolkit for generating analytic specifications for data visualization from natural language queries. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):369–379.
- OpenAI. 2024. Learning to reason with llms.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Paul Parsons. 2021. Understanding data visualization design practice. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):665–675.
- Ghulam Jilani Quadri, Arran Zeyu Wang, Zhehao Wang, Jennifer Adorno, Paul Rosen, and Danielle Albers Szafir. 2024. Do you see what i see? a qualitative study eliciting high-level visualization comprehension. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 1–26.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR.
- Donghao Ren, Matthew Brehmer, Bongshin Lee, Tobias Höllerer, and Eun Kyoung Choe. 2017. ChartAccent: Annotation for data-driven storytelling. In *2017 IEEE Pacific Visualization Symposium (PacificVis)*, pages 230–239. Ieee.
- Theresa-Marie Rhyne. 2017. Applying color theory to digital media and visualization. In *Proceedings of the 2017 CHI conference extended abstracts on human factors in computing systems*, pages 1264–1267.
- María Teresa Rodríguez, Sérgio Nunes, and Tiago Devezas. 2015. Telling stories with data visualization. In *Proceedings of the 2015 Workshop on Narrative & Hypertext*, pages 7–11.
- Ruth Rosenholtz, Yuanzhen Li, Jonathan Mansfield, and Zhenlan Jin. 2005. Feature congestion: a measure of display clutter. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 761–770.
- Baptiste Roziere, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Romain Sauvestre, Tal Remez, et al. 2023. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950*.
- Subham Sah, Rishab Mitra, Arpit Narechania, Alex Endert, John Stasko, and Wenwen Dou. 2024. [Generating analytic specifications for data visualization from natural language queries using large language models](#). *Preprint*, arXiv:2408.13391.
- Tobias Skog, Sara Ljungblad, and Lars Erik Holmquist. 2003. Between aesthetics and utility: designing ambient information visualizations. In *IEEE Symposium on Information Visualization 2003 (IEEE Cat. No. 03TH8714)*, pages 233–240. IEEE.
- Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. 2024. Scaling llm test-time compute optimally can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*.
- Chase Stokes, Cindy Xiong Bearfield, and Marti A Hearst. 2023. The role of text in visualizations: How annotations shape perceptions of bias and influence predictions. *IEEE Transactions on Visualization and Computer Graphics*.
- Chase Stokes, Vidya Setlur, Bridget Cogley, Arvind Satyanarayan, and Marti A Hearst. 2022. Striking a balance: Reader takeaways and preferences when integrating text and charts. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):1233–1243.

- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.
- Martijn Tennekes and Edwin de Jonge. 2014. Tree colors: color schemes for tree-structured data. *IEEE transactions on visualization and computer graphics*, 20(12):2072–2081.
- Yuan Tian, Weiwei Cui, Dazhen Deng, Xinjing Yi, Yurun Yang, Haidong Zhang, and Yingcai Wu. 2024. [Chartgpt: Leveraging llms to generate charts from abstract natural language](#). *IEEE Transactions on Visualization and Computer Graphics*, page 1–15.
- Jonathan Uesato, Nate Kushman, Ramana Kumar, Francis Song, Noah Siegel, Lisa Wang, Antonia Creswell, Geoffrey Irving, and Irina Higgins. 2022. Solving math word problems with process-and outcome-based feedback. *arXiv preprint arXiv:2211.14275*.
- Pere-Pau Vázquez. 2024. Are llms ready for visualization? *arXiv preprint arXiv:2403.06158*.
- Ante Wang, Linfeng Song, Ye Tian, Baolin Peng, Dian Yu, Haitao Mi, Jinsong Su, and Dong Yu. 2024a. Litesearch: Efficacious tree search for llm. *arXiv preprint arXiv:2407.00320*.
- Chaojie Wang, Yanchen Deng, Zhiyi Lv, Shuicheng Yan, and An Bo. 2024b. Q*: Improving multi-step reasoning for llms with deliberative planning. *arXiv preprint arXiv:2406.14283*.
- Xingyao Wang, Zihan Wang, Jiateng Liu, Yangyi Chen, Lifan Yuan, Hao Peng, and Heng Ji. 2024c. [MINT: Evaluating LLMs in multi-turn interaction with tools and language feedback](#). In *The Twelfth International Conference on Learning Representations*.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. [Self-consistency improves chain of thought reasoning in language models](#). In *The Eleventh International Conference on Learning Representations*.
- Ze Zhong Wang, Lovisa Sundin, Dave Murray-Rust, and Benjamin Bach. 2020. Cheat sheets for data visualization techniques. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–13.
- Chengyue Wu, Yixiao Ge, Qiushan Guo, Jiahao Wang, Zhixuan Liang, Zeyu Lu, Ying Shan, and Ping Luo. 2024. Plot2code: A comprehensive benchmark for evaluating multi-modal large language models in code generation from scientific plots. *arXiv preprint arXiv:2405.07990*.
- Renqiu Xia, Bo Zhang, Hancheng Ye, Xiangchao Yan, Qi Liu, Hongbin Zhou, Zijun Chen, Min Dou, Botian Shi, Junchi Yan, et al. 2024. Chartx & chartvlm: A versatile benchmark and foundation model for complicated chart reasoning. *arXiv preprint arXiv:2402.12185*.
- Anran Xu, Shitao Fang, Huan Yang, Simo Hosio, and Koji Yatani. 2024a. [Examining human perception of generative content replacement in image privacy protection](#). In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI '24, New York, NY, USA. Association for Computing Machinery.
- Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao Dong. 2024b. Imagereward: Learning and evaluating human preferences for text-to-image generation. *Advances in Neural Information Processing Systems*, 36.
- Zhengzhuo Xu, Sinan Du, Yiyan Qi, Chengjin Xu, Chun Yuan, and Jian Guo. 2023. Chartbench: A benchmark for complex visual reasoning in charts. *arXiv preprint arXiv:2312.15915*.
- Fumeng Yang, Yuxin Ma, Lane Harrison, James Tompkin, and David H. Laidlaw. 2023. [How can deep neural networks aid visualization perception research? three studies on correlation judgments in scatterplots](#). In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, New York, NY, USA. Association for Computing Machinery.
- Zhiyu Yang, Zihan Zhou, Shuo Wang, Xin Cong, Xu Han, Yukun Yan, Zhenghao Liu, Zhixing Tan, Pengyuan Liu, Dong Yu, et al. 2024. Matplotagent: Method and evaluation for llm-based agentic scientific data visualization. *arXiv preprint arXiv:2402.11453*.
- Fatemeh Pesaran Zadeh, Juyeon Kim, Jin-Hwa Kim, and Gunhee Kim. 2024. Text2chart31: Instruction tuning for chart generation with automatic feedback. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 11459–11480.
- Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex Smola. 2023. [Automatic chain of thought prompting in large language models](#). In *The Eleventh International Conference on Learning Representations*.
- Jian Zhao, Shenyu Xu, Senthil Chandrasegaran, Chris Bryan, Fan Du, Aditi Mishra, Xin Qian, Yiran Li, and Kwan-Liu Ma. 2021. Chartstory: Automated partitioning, layout, and captioning of charts into comic-style narratives. *IEEE transactions on visualization and computer graphics*, 29(2):1384–1399.

Contents

1	Introduction	1
1.1	Related Work	2
2	C²: CHARTAF	3
2.1	Towards High-performing Feedback	3
2.2	CHARTAF-S (\tilde{f}_{AF})	3
2.3	CHARTAF-G (f_{AF})	4
3	C²: CHARTUIE-8K	4
4	Empirical Study	5
4.1	Preliminary and Notations	5
4.2	Test-time Scaling with CHARTAF	5
4.3	In-context Tuning with CHARTAF	5
4.4	CHARTUIE-8K Experiments	7
5	Discussion and Impact of C²	7
5.1	CHARTAF Enables Scalable Feed- back	7
5.2	Unlocking Large-scale Data with C ²	8
6	Limitations	9
7	Human Study Ethical Consideration	9
A	CHARTAF Details	15
A.1	Basic Criteria References	15
A.2	CHARTAF Pseudocode	15
B	CHARTUIE-8K Details	24
B.1	CHARTUIE-8K Pseudocode	24
B.2	Example User Study Initial Instruc- tions	31
C	Chart Generation User Query Examples	33
D	Human Study Details	35
D.1	Compensation and Qualification	35
D.2	Sanity Check	35
D.3	CHARTAF Human Study	35
D.4	CHARTUIE-8K Human Study	36
E	LLM Configurations	41
F	Code Error Rate	41
G	Icon Attribution	41

A CHARTAF Details

A.1 Basic Criteria References

See Table 4.

Criteria Category	References
Chart Type	Figueiras (2013); Jung et al. (2017); Islam and Jin (2019b); Midway (2020)
Visual Embellishment	Bateman et al. (2010); Borgo et al. (2012); Andry et al. (2021)
Text	Chi et al. (2015); Stokes et al. (2022, 2023)
Color	Healey (1996); Lee et al. (2012); Tennekes and de Jonge (2014); Rhyne (2017)
Annotation	Lee et al. (2015); Ren et al. (2017); Lavrič et al. (2017); Chen et al. (2018); Kim et al. (2021); Zhao et al. (2021)
Aesthetics	Skog et al. (2003); Cawthon and Moere (2007); Kim and Park (2013); Harrison et al. (2015)
Visual Clutter	Rosenholtz et al. (2005); Ellis and Dix (2007); Ajani et al. (2022)

Table 4: References that form the theoretical background of the Basic Criteria

A.2 CHARTAF Pseudocode

We present the pseudocode of CHARTAF in Alg. 1, along with the prompts used within it. These prompts can be accessed by clicking on the highlighted phrases in the pseudocode.

We first introduce the notations. Let $inst$ represent the initial instruction, qst the follow-up questions, and ans the answers to qst . Let d represent the dataset, and d_{attr} its attributes. f_{AF} denotes the backbone LLM for feedback generation. Let $mode$ be one of two values: "Scalar" or "Granular." Finally, let $code_{gen}$ be the code that generates a chart. This code is produced by the chart-generating LLM, using the prompt *Generate*. The arguments in the prompt— $data_path$, $data$, $file_index$, $initial_instruction$, $questions$, and $answers$ —should be set to the data path for the dataset, d , the index for the resulting image, $inst$, qst , and ans , respectively. Denote img as the chart generated by executing $code_{gen}$. The procedure returns either a single scalar value, s_N , or fine-grained natural language feedback, s_N , depending on the value of $mode$.

Algorithm 1 ChartAF

```
1: procedure CHARTAF( $inst, qst, ans, d_a, d, f_{AF}, mode, code_{gen}, img$ )
2:    $p_{tpa} \leftarrow TPA_{AF}$ 
3:    $p_{tpa}.format($ 
4:      $initial\_instruction:=inst,$ 
5:      $tasks:=task,$ 
6:      $data:=d$ 
7:   )
8:    $\langle tsk, prps, aud \rangle \leftarrow f_{AF}(p_{tpa})$ 
9:    $p_{crt} \leftarrow Criteria$ 
10:   $p_{crt}.format($ 
11:     $initial\_instruction:=inst,$ 
12:     $questions:=qst,$ 
13:     $answers:=ans,$ 
14:     $tasks:=tsk,$ 
15:     $purpose:=prps,$ 
16:     $audience:=aud$ 
17:  )
18:   $crt \leftarrow f_{AF}(p_{crt})$ 
19:   $p_{qst} \leftarrow Criteria_Q$ 
```

Algorithm 1 ChartAF (Continued)

```
20:   $p_{qst}$ .format(  
21:    task:= $tsk$ ,  
22:    purpose:= $prps$ ,  
23:    audience:= $aud$ ,  
24:    criteria:= $crt$   
25:  )  
26:   $crt_{qst} \leftarrow f_{AF}(p_{qst})$   
27:   $p_{eval} \leftarrow Evaluate$   
28:   $p_{eval}$ .format(  
29:    evaluation_questions:= $crt_{qst}$ ,  
30:    initial_instruction:= $inst$   
31:  )  
32:   $s \leftarrow f_{AF}(img, p_{eval})$   
33:   $s_{\mathbb{N}} \leftarrow$  Ratio of "yes" responses in  $s$   
34:  if  $mode$  is "Scalar" then  
35:    return  $s_{\mathbb{N}}$   
36:  end if  
37:   $s_{\mathcal{N}} \leftarrow$  Feedback in  $s$   
38:   $\langle rtn, dsc, edt, add \rangle \leftarrow$  Classification of  
39:    feedback according to the tag in  $s_{\mathcal{N}}$   
40:   $p_{cf} \leftarrow CF$   
41:   $p_{cf}$ .format(  
42:    initial_instruction:= $inst$ ,  
43:    code:= $code_{gen}$ ,  
44:    attributes:= $d_{attr}$ ,  
45:    retain:= $rtn$ ,  
46:    discard:= $dsc$ ,  
47:    edit:= $edt$ ,  
48:    add:= $add$   
49:  )  
50:   $s_{\mathcal{N}} \leftarrow f_{AF}(p_{cf})$   
51:  return  $s_{\mathcal{N}}$   
52: end procedure
```

Generating a Chart (*Generate*)

You are an expert data visualizer.

The following instruction asks you to generate code for data visualization of the underlying data file that we have attached. I will give you the data, but you can ignore some parts from the data if it is not necessary and unrelated to the instruction. Assume that the data file that has been attached in the path "{data_path}" in the generated code. The file format of the data is "**f".json or ".csv".** Your code should include loading the data file, and check and verify the data type and representation of the data to avoid errors while executing.

```
<start of data format>  
{data}  
<end of data format>
```

Your code should also automatically download the final visualization in a lower level directory (contained within the current directory) named "plots_d2c". You **MUST** name your final generated

visualization as "{file_index}.png". You can freely choose package(s) that work best to make the visualization.

Here is the instruction set:

```
<start of initial instruction>
{initial_instruction}
<end of initial instruction>
```

```
<start of further instruction>
Questions:
{questions}
Answers:
{answers}
<end of further instruction>
```

Ensure you use this code format in order to avoid errors, and only give the executable Python Code.

CHARTAF Task Purpose and Audience Inference (TPA_{AF})

You are a data visualization expert. Given the data and user request, your task is to analyze the user request to (1) select the most suitable task that the user is expecting from the list of various tasks in data visualization, (2) specifically figure out the purpose of the user's request in data visualization, and (3) prospective audience of the data visualization.

```
<start of user request>
{initial_instruction}
<end of user request>
```

```
<start of data format>
{data}
<end of data format>
```

```
<start of various task types>
{tasks}
<end of various task types>
```

We use data visualization tasks presented in [Choe et al. \(2017\)](#).

List of Tasks (*task*)

- Show External Context
Uncaptured data provided by the self-tracker to understand and explain a phenomenon shown in the data.
- Show confirmation
Collected data confirms existing knowledge.
- Show Contradiction
Collected data contradicts existing knowledge.

- **Focus on Identifying value**
Explicitly specify the measured value, its range for one or more clearly identified data points, or the difference between two measured values.
- **Focus on Identifying extreme**
Explicitly state the identities of the data points possessing extreme values of the measure variable.
- **Focus on Identifying reference**
Explicitly state the values of categorical variables, labels from the axes, or legends.
- **Comparison by Time Segmentation**
Compare measured values segmented by time.
- **Comparison by Multiple services**
Compare the same data type from two or more services.
- **Comparison against external data**
Bringing in external data for comparison.
- **Comparison by Factor**
Compare measured values by a factor (other than time).
- **Comparison by Instances**
Compare two specific instances.
- **Show Trend**
Describe changes over time.
- **Value judgement**
Convey positive or negative connotations about the data.
- **Distribution with variability**
Explicitly state the variability of measured values.
- **Distribution By Category**
Explicitly describe the variation of measured values across all or most of the values of a categorical variable.
- **Correlation**
Specify the direct relationship between two variables (but not as comparison).
- **Outlier**
Explicitly point out outliers or state the effect of outliers.
- **Sumarization of data**
Summary of collected data (such as number of data points, duration of tracking, and averages).
- **Prediction/Forecasting**
Predict the future based on the collected data.

CHARTAF Criteria Establishment (*Criteria*)

You are a data visualization expert. You are given basic essential requirements of chart, user instruction, user request with QA, tasks that must be covered by the chart, purpose of the chart, and prospective audience of the chart.

Your task is to develop a personalized, detailed, and objective list of criteria, building on the

basic criteria, to evaluate a data visualization (chart). These criteria should be based on the user instruction, the user request through questions and answers, tasks at hand, intended purpose, and the prospective audience.

<start of basic criteria>

- **Chart Type**
Choose a chart type that aligns with the given purpose, task, and audience. The chart type should effectively convey the intended message; for example, bar charts are ideal for comparing quantities for limited number of categorical data, while line charts show trends over time. The choice must consider the inherent spacing requirements and the context in which the chart will be used, ensuring clarity and comprehension.
- **Visual Embellishment**
Use embellishments to enhance understanding without overwhelming the data. Visual embellishments, like icons, patterns, or textures, should be used sparingly and purposefully to make the chart memorable and engaging while maintaining a balance that does not distract from the core data.
- **Text**
Prioritize legibility and adhere to consistent textual criteria. Text elements, such as legends, titles, and labels, should be legible and easy to read, with sufficient contrast against the background. Consistent font size, style, and placement should be maintained to create a cohesive visual narrative that guides the audience's understanding.
- **Color**
Use color purposefully and sparingly to convey meaning. Choose a limited palette that enhances readability and highlights key data points, considering color statistics and opponent processing principles (contrasting colors for clarity). This helps ensure accessibility for viewers with color vision deficiencies.
- **Annotation**
Emphasize critical data while minimizing irrelevant details. Use annotations strategically to draw attention to important insights, trends, or outliers, and smooth over or de-emphasize less significant data points, ensuring the chart communicates its key message effectively.
- **Aesthetics**
Tailor aesthetics to the chart's purpose, audience, and context. Consider the chart's purpose, the target audience, and the presentation environment when designing aesthetics, including compact spacing and visual hierarchy. This ensures the chart is both functional and appealing, maximizing its impact and effectiveness
- **Visual Clutter**
Optimize the chart size to fit its content and context, balancing data and available space to prevent clutter or excessive white space while maintaining readability. Manage visual elements by minimizing overcrowding and overlapping, adequately spacing text, data points, and annotations, removing unnecessary details, and maintaining a clean layout to enhance clarity. Segmentation of complex charts or data visualizations can also be employed if the visual complexity is high, breaking down the data into smaller, more manageable parts for easier interpretation. It is important to emphasize key data by using size, color, and opacity to highlight critical insights while downplaying less relevant information for a focused presentation.

<end of basic criteria>

<start of user instruction>
{initial_instruction}
<end of user instruction>

<start of user request through QA>
{questions}
{answers}
<end of user request through QA>

<start of tasks>
{tasks}
<end of tasks>

<start of purpose>
{purpose}
<end of purpose>

<start of prospective audience>
{audience}
<end of prospective audience>

Note that the interactivity of the chart, file format, credibility and integrity of data source, summary statistics do not need to be considered. The quality of the data visualization to general audience is the only subject to be considered.

Think about the essential chart component requirements that align with the task, purpose, and user request.

CHARTAF Generate Criteria Questions (*Criteria_Q*)

You are an expert critic. You will be given wanted tasks, intended purpose, prospective audience, and established criteria for the chart that you gave in the previous prompt. Your task is to create a list of Yes/No questions that checks if the generated chart satisfies the established criteria. Use the established criteria as a reference, but avoid applying them directly when crafting questions to evaluate the chart.

<start of task>
{task}
<end of task>

<start of purpose>
{purpose}
<end of purpose>

<start of prospective audience>
{audience}
<end of prospective audience>

<start of established criteria>
{criteria}
<end of established criteria>

"Yes" should be treated as satisfaction, while "No" should be a dissatisfaction.

Here is a detailed protocol for making questions:

First, create questions according to the criteria, tasks, purpose, and audience. Extra questions that the criteria do not cover can be generated, yet it must help judge evaluating the chart. Lastly, summarize similar questions and rank the questions so that the first question is the most important and the last question is the least important.

Your output should follow the format below:

””

Question 1 : [Question]

Question 2 : [Question]

...

””

CHARTAF Evaluation (*Evaluate*)

You are an expert evaluator (judge, critic) of the attached data visualization image.

<start of evaluation questions>

{evaluation_questions}

<end of evaluation questions>

The evaluation questions consist of YES/NO questions; the answer for each question **MUST** be either YES or NO. Don't give anything else like N/A. With the answers, you need to give feedback. When answering the questions, follow the step-by-step protocol below:

1. Determine and tag whether the question is subjective or fact-checking
 - Fact-checking
Verify if the chart image meets the criteria directly based on the visual content. If the image shows any deviation from the criteria, answer NO. If the image meets the criteria, answer YES.
 - Subjective
Consider whether the image meets the criteria based on visual appeal, clarity, and other subjective measures. Provide reasons for both YES and NO answers. If there is clear evidence to support a YES and no substantial reasons to support a NO, answer YES. Answer NO otherwise.
2. Answer the questions and provide feedback
After answering each question, provide feedback explaining your evaluation. List potential improvements categorized as RETAIN, DISCARD, EDIT, or ADD if necessary.
Feedback Classification:
 - RETAIN
Identify and specify any elements that should be retained even after the improvement.
 - DISCARD
Identify and specify any elements that should be discarded for better visualization.
 - EDIT
Specify edits needed in the image to satisfy the user's request. Provide examples if applicable.
 - ADD
Identify and specify elements that should be added for better visualization of the user's initial prompt.

To help your task, here is the user's initial prompt.

```
<start of initial prompt>  
{initial_instruction}  
<end of initial prompt>
```

CHARTAF Generate Code Feedback (CF)

You are an expert software engineer on the Quality Assurance team Your task is to provide feedback on the code based on the critic's feedback on the result of the code. The code's goal is to successfully draw a chart, fulfilling the user's needs. You will be given the user's needs, the original code, the critic's feedback, the data attributes, and the resulting image of the code.

Here is the user's needs.

```
<start of needs>  
{initial_instruction}  
<end of needs>
```

Here is the original code.

```
<start of the code>  
{code}  
<end of the code>
```

Here are the data attributes.

```
<start of the attributes>  
{attributes}  
<end of the attributes>
```

Here is the critic's feedback.

```
<start of feedback>  
Elements to RETAIN  
—  
{retain}
```

Elements to DISCARD

```
—  
{discard}
```

Elements to EDIT

```
—  
{edit}
```

Elements to ADD

```
—  
{add}  
<end of feedback>
```

Your task is to provide feedback on the code for debugging and offering better data visualization. Specifically, focus on cases where the image does not correctly reflect the intended output, even though the code appears correct. Follow these steps:

1. Review the Evaluation Feedback

Examine the feedback, especially noting where the image does not align with the expected results despite the code being correct.

2. Analyze the Feedback

Determine what changes are necessary in the code to correct errors and enhance the output based on the feedback. If there are potential errors that may occur, feel free to provide feedback on those lines. Again, your task is not only to offer better data visualization but also to debug the code.

3. List your feedback on the code, and make sure such modifications help generate the executable code.

Explain the modification, log the lines of code that should be modified, and log lines of new code that can be implemented. When logging the code, log the line number as well, where the original code lies, and where the new code should be put.

B CHARTUIE-8K Details

B.1 CHARTUIE-8K Pseudocode

We present the pseudocode of CHARTUIE-8K in Alg. 2, along with the prompts used within it. These prompts can be accessed by clicking on the highlighted phrases in the pseudocode.

We first introduce the notations. Let d represent the underlying dataset. If it exists, let d_{ttl} denote the title of d ; otherwise, set d_{ttl} to "unknown." Similarly, if the topic of d is provided, let d_{tpc} represent it. Finally, let f_{uie} refer to the backbone LLM for CHARTUIE-8K. Here, $f_{uie} := \text{GPT-4o}$.

The procedure returns a list of queries, Q , generated by the algorithm. Each tuple in Q includes the initial instruction ($inst$), column labels for visualization (data attributes, d_{attr}), a selected task (tsk), a visualization purpose ($prps$), a text description of the target audience (aud), follow-up questions (qst), and answers (ans). The model f_{uie} infers d_{attr} , tsk , $prps$, and aud from $inst$ and follow-up questions to clarify user preferences, of which only $q\%$ are answered.

Algorithm 2 CHARTUIE-8K

```
1: procedure UIE-8K( $d, d_{ttl}, d_{tpc}, f_{uie}$ )
2:    $p_{ct} \leftarrow \text{Select}_{ct}$ 
3:    $p_{ct}.\text{format}(\text{data}:=d)$ 
4:    $C_{type} \leftarrow f_{uie}(p_{ct})$ 
5:    $p_{annot} \leftarrow \text{Select}_{annot}$ 
6:    $p_{annot}.\text{format}(\text{data}:=d,$ 
7:      $\text{data\_title}:=d_{ttl},$ 
8:      $\text{topic}:=d_{tpc}$ 
9:   )
10:   $annot \leftarrow f_{uie}(p_{annot})$ 
11:   $j_{max} \leftarrow \min\{15, \text{the length of } C_{type}\}$ 
12:   $Q \leftarrow []$ 
13:  for  $j \in \{0, 1, \dots, j_{max} - 1\}$  do
14:    for  $word \in \{50, 100\}$  do
15:      if  $j \leq 2$  then
16:         $p_{trg} \leftarrow \text{Trigger}_{annot}$ 
17:         $p_{trg}.\text{format}(\text{data}:=d,$ 
18:           $\text{annotations}:=annot,$ 
19:           $\text{data\_title}:=d_{ttl},$ 
20:           $\text{topic}:=d_{tpc},$ 
21:           $\text{chart\_type}:=C_{type}[j]$ 
22:           $\text{word\_count}:=word$ 
23:        )
24:       $Q \leftarrow Q \cup \{p_{trg}\}$ 
25:  return  $Q$ 
```

Algorithm 2 CHARTUIE-8K (Continued)

```
26:         else
27:              $p_{trg} \leftarrow Trigger_{annot}$ 
28:              $p_{trg}.format($ 
29:                  $data:=d,$ 
30:                  $data\_title:=d_{ttl},$ 
31:                  $topic:=d_{tpc},$ 
32:                  $chart\_type:=C_{type}[j]$ 
33:                  $word\_count:=word$ 
34:             )
35:         end if
36:          $inst \leftarrow f_{uie}(p_{trg})$ 
37:          $p_{tpa} \leftarrow TPA_{uie}$ 
38:          $p_{tpa}.format($ 
39:              $initial\_instruction:=inst,$ 
40:              $data:=d,$ 
41:              $task:=task$ 
42:         )
43:          $\langle d_a, tsk, prps, aud \rangle \leftarrow f_{uie}(p_{tpa})$ 
44:          $p_{qst} \leftarrow Question_{uie}$ 
45:          $p_{qst}.format($ 
46:              $data:=d,$ 
47:              $initial\_instruction:=inst,$ 
48:              $attributes:=d_{attr},$ 
49:              $audience:=aud,$ 
50:              $tasks:=tsk,$ 
51:              $purpose:=prps$ 
52:         )
53:          $qst \leftarrow f_{uie}(p_{qst})$ 
54:          $p_{ans} \leftarrow Answer_{uie}$ 
55:         if  $word == 50$  then
56:              $q\% \leftarrow 25$ 
57:         else
58:              $q\% \leftarrow 50$ 
59:         end if
60:          $p_{ans}.format($ 
61:              $initial\_instruction:=inst,$ 
62:              $q\_percent:=q\%,$ 
63:              $purpose:=prps,$ 
64:              $f\_response:= qst$ 
65:         )
66:          $ans \leftarrow f_{uie}(p_{ans})$ 
67:          $Q.append($ 
68:              $\langle$ 
69:                  $inst, d_a, tsk, prps,$ 
70:                  $aud, qst, ans$ 
71:              $\rangle$ 
72:         )
73:     end for
74: end for
75: return  $Q$ 
76: end procedure
```

We use chart types presented in Hess (2022).

Chart Type Selection (*Select_{ct}*)

<start of data example format>

{data}

<end of data example format>

Look at the following list and select one or more chart types appropriate for the data. Try to choose as many different charts as possible. Consider the purpose and chart type that can visualize the data well.

But Chart types that are not reasonable for data visualization of the data example format I attached must be excluded. Respond with the chart types that are compatible with the data. Please include only one chart type that is similar to each other. Comparison, Correlation, Part-to-whole & hierarchical, Data over time (temporal), Distribution, Geospatial & other are the purposes, and the chart types below are suitable for the above purpose.

Look closely at the characteristics of the data, and the annotation should be one that produces as little clutter as possible.

<start of the chart type list>

1. Comparison

- Bar chart
- Column chart
- Grouped bar/column chart
- Lollipop chart
- Bullet chart
- Dot plot
- Dumbbell
- Pictogram
- Icon chart
- Range chart
- Radial bar chart
- Parallel coordinates
- Radar chart
- Nightingale chart
- Waterfall chart
- Matrix chart
- Small multiples
- Word cloud
- Slope chart
- Table chart
- Categorical scatter plot
- Quadrant chart

2. Correlation

- Heatmap
- Bubble chart
- Scatter plot

- Connected scatter plot
- Hexagonal binning
- Contour plot

3. Part-to-whole & hierarchical

- Stacked bar/column chart
- Diverging bar chart
- Population pyramid
- Icon array
- Waffle chart
- Pie chart
- Donut chart
- Semi-circle donut chart
- Marimekko chart
- Treemap
- Circular treemap
- Convex treemap
- Dendrogram
- Venn diagram
- Euler diagram
- Circular gauge
- Sunburst chart
- Funnel & pyramid chart

4. Data over time (temporal)

- Area chart
- Stacked area chart
- Stream graph
- Bump chart
- Bump area chart
- Line chart
- Spline chart
- Step line chart
- Candlestick chart
- Gantt chart
- Barcode chart
- OHLC chart

5. Distribution

- Density plot
- Ridgeline plot
- Horizon chart
- Histogram
- Radial histogram
- Strip plot
- Jitter plot
- One-dimensional heatmap

- Beeswarm chart
- Box chart
- Violin plot

6. Geospatial & other

- Geographic heatmap
- Choropleth map
- Tile map
- Chord diagram
- Arc diagram
- Sankey
- Network diagram
- Flowchart

<end of the chart type list>

Your response should ONLY contain the chart types. Do not include anything else.

Annotation Selection (*Select_{annot}*)

<start of data example format>

{data}

<end of data example format>

<start of data details format>

{data_title}

{topic}

<end of data details format>

Look at the following list and select one or more annotations appropriate for the data. Choose two annotations.

<start of the annotation list>

1. Text Annotations:

Description: Data-driven text annotations display values linked to chart elements, such as data points in a scatterplot. They draw attention to specific elements by highlighting their values. Purpose: When only some elements are annotated, the intent is to focus the viewer's attention on those before examining others. Other Uses: Non-data-driven annotations can provide context, orientation, or editorial comments.

2. Shapes:

Description: Shape annotations include lines, arrows, rectangles, and other shapes. They can highlight or enclose specific chart elements to emphasize or compare them. Data-Driven Use: Some shapes, like trend lines, are calculated from the underlying data.

3. Highlights:

Description: Highlights modify the appearance of chart elements (e.g., size, color) to emphasize or reduce their importance. Purpose: Used to distinguish certain elements from others, making them stand out visually

<end of the annotation list>

Look closely at the characteristics of the data, and the annotation should be one that produces as little clutter as possible. Also, refer to the data details to create as practical and realistic instructions as possible. Your response should ONLY contain the annotations, description, purpose, and uses. Do not include anything else.

Emulating with annotations (*Trigger_{annot}*)

You are an expert user emulator.

<start of data example format>
{data}
<end of data example format>

<start of annotations format>
{annotations}
<end of annotations format>

<start of data details format>
{data_title}
{topic}
<end of data details format>

Given a data format, imagine a chart that visualizes this data as the final output you want from the service provider. It MUST be a chart that can be created using only data columns. Consider what purpose the data has and the practical purpose of visualization and include it in the instructions. You need to imagine a chart with {chart_type} and given annotations that utilizes the data format. If there are multiple given data formats, imagine a chart with {chart_type} and given annotations that utilizes all the data formats. Since you are an amateur user, your instruction will be partially SUBJECTIVE and NOT DETAILED. Also, refer to the data details to create as practical and realistic instructions as possible. Instructions must reflect the context of the data. To emulate a real-world user your instruction should be {word_count} in size (word count). Do not include data path in the instruction. Your response should ONLY contain the user emulated instruction. Do not include anything else.

Emulating without annotations (*Trigger_{annot'}*)

You are an expert user emulator.

<start of data example format>
{data}
<end of data example format>

<start of data details format>
{data_title}
{topic}
<end of data details format>

Given a data format, imagine a chart that visualizes this data as the final output you want from the service provider. It MUST be a chart that can be created using only data columns. You need to imagine a chart with {chart_type} that utilizes the data format. If there are multiple given data formats, imagine a chart with {chart_type} that utilizes all the data formats. Since you are an

amateur user, your instruction will be partially SUBJECTIVE and NOT DETAILED. Also, refer to the data details to create as practical and realistic instructions as possible. Instructions must reflect the context of the data. To emulate a real-world user your instruction should be {word_count} in size (word count). Your response should ONLY contain the user emulated instruction. Do not include anything else.

Task Purpose and Audience Inference for UIE (TPA_{vie})

You are a data visualization expert. Given the data and user request, your task is to analyze the user request to figure out the (1) data attributes needed for data visualization, (2) select the most suitable task that the user is expecting from the list of various tasks in data visualization, (3) specifically figure out the purpose of the user's request in data visualization, and (4) prospective audience of the data visualization.

<start of user request>
{initial_instruction}
<end of user request>

<start of data format>
{data}
<end of data format>

<start of various task types>
{task}
<end of various task types>

For data attributes needed in data visualization, store them in query['Data Attribute']. Data attributes MUST match exactly the column names of the data. Store the selected task from the task types in query['Task']. Store the purpose of visualization in query['Purpose']. Store the prospective audience in query['Audience'].

Please reply in the same format without altering the key value.

{"Data Attribute": None, "Task": None, "Purpose": None, "Audience": None}

But, please make sure there is no ' in each keys and values. Use only " for the response. But when you write a value sentence or each data attribute's title, you only can use '. Unless you are writing a sentence or each data attribute's title, you should never include ' in response. If there are multiple pieces of data, there is no need to reveal which file each is. Please consider that JSON conversion must be done properly.

List Preference Questions ($Question_{vie}$)

You are an expert data visualization analyst. Given data, data attribute, tasks, prospective audience, purposes of the chart(data visualization) from the initial instruction (user request) of the user, you have a 2-step task to do.

(1) First, your task is to figure out the essential chart attribute requirements that the chart must have in order to satisfy such tasks and purposes (2) Then, create a list of questions to the user if the user have specific chart attribute preference for effective data visualization.

Do NOT include anything else in your response other than the list of questions. Your questions should be primarily focused on retrieving the user's preferences. Do NOT include any questions related to (1) interactivity of the chart, and (2) the file format of the chart.

<start of data format>

```
{data}
<end of data format>

<start of data visualization instruction>
{initial_instruction}
<end of data visualization instruction>

<start of data attributes>
{attributes}
<end of data attributes>

<start of prospective audience>
{audience}
<end of prospective audience>

<start of type of tasks>
{tasks}
<end of type of tasks>

<start of purpose>
{purpose}
<end of purpose>
```

Answer Preference Questions (*Answer_{uie}*)

You are an expert user emulator.

```
<start of user request>
{initial_instruction}
<end of user request>
```

Imagine the best result that can be achieved based on the given instructions and is the final output you want from the service provider with your purpose given below. Based on your imagination you are to respond to the service provider. Since you are an amateur user you should concretely answer {q_percent}% of the questions you think are most important, and the remaining questions should be answer with uncertainty, e.g. "I am not sure", "I do not know", "I have no specific preference". Finally, remember that you are looking for an image, not an interactive data visualization.

```
<start of purpose>
{purpose}
<end of purpose>
```

```
<start service provider message>
{f_response}
<end service provider message>
```

Also, remember that the service provider can not see the image you have access to. Your response should ONLY contain the user emulated response. Do not include anything else.

B.2 Example User Study Initial Instructions

We provide user study initial instruction examples curated by the LLM user participants below.

User Initial Instruction Example 1

Please make a bar chart of R&D investment (in billions of USD) according to country. Make each bar green, except for Japan and India, which you should make blue. Title the chart, "R&D Investment by Country." Include a red line that connects the maximum of each bar. And add a blue dash line where the average is across the whole chart (including the average value in Billions out to two decimal places).

User Initial Instruction Example 2

Chart grades over 12 months. Have different colors for each student and label by names. Y axis will be percentage and x axis months. Title will be Students Grades Over the Year.

C Chart Generation User Query Examples

Plot2Code Example

The figure generated by the provided Python code consists of four subplots, each displaying a different transformation of the same image. The image is created using a mathematical function that generates a grid of values.

The first subplot shows the image rotated by 30 degrees. The second subplot shows the image skewed by 30 degrees along the x-axis and 15 degrees along the y-axis. The third subplot shows the image reflected along the x-axis and scaled by 0.5 along the y-axis. The fourth subplot combines all these transformations: rotation, skewing, scaling, and translation.

Each subplot has a yellow dashed line indicating the intended extent of the image. The x-axis limits are set from -5 to 5, and the y-axis limits are set from -4 to 4 for all subplots.

To recreate this figure, you would need to generate the same image and apply the same transformations in each subplot. The transformations are applied using an affine transformation matrix, which allows for rotation, skewing, scaling, and translation. The specific parameters for each transformation are mentioned above.

The figure is displayed using the matplotlib library in Python, which would need to be installed and imported to recreate the figure. The numpy library is also used to generate the image and would need to be installed and imported as well.

The image is generated using a mathematical function that creates a grid of values between -3.0 and 3.0 with a step of 0.25. This function uses the numpy exp function to calculate the exponential of the negative square of each value in the grid. The result is a 2D array of values, which is used as the image in the subplots.

The transformations applied to the image in each subplot are done using the Affine2D function from the matplotlib.transforms module. This function creates a 2D affine transformation matrix that can be used to apply various transformations to the image. The specific transformations and their parameters are as follows:

- Rotation: 30 degrees
- Skew: 30 degrees along the x-axis and 15 degrees along the y-axis
- Scale: -1 along the x-axis (which reflects the image) and 0.5 along the y-axis
- Translation: 0.5 along the x-axis and -1 along the y-axis

The extent of the image in each subplot is set to [-2, 4, -3, 2], and the image is clipped to these bounds. The yellow dashed line indicating the intended extent of the image is drawn using the plot function with these bounds.

Finally, the figure is displayed using the show function from the matplotlib.pyplot module.

MatPlotBench's Simple Instruction Example

Create a Python script using matplotlib to generate a specific plot with the following detailed parameters:

Initialize a figure with a custom size of 7.5 by 7.5 inches. Add a single axis to the figure with a custom aspect ratio and specified position. Define X as a linear space from 0.5 to 3.5 with 100 elements. Calculate Y1 as 3 plus the cosine of X, Y2 as 1 plus the cosine of $1+X/0.75$ divided by 2, and Y3 as random values uniformly distributed between Y1 and Y2.

Set major and minor locators for both x and y axes with major intervals of 1 and minor intervals of 4. Set minor formatter for the x-axis to display values with two decimal places. Limit the x and y axes to a range from 0 to 4.

For major ticks, set the width to 1.0, length to 10, and label size to 14. For minor ticks, set the width to 1.0, length to 5, label size to 10, and label color to '.25'. Add a grid with these specific attributes:

- Linestyle set to "-" (dashed).
- Linewidth of 0.5.
- Color set to '.25' (a shade of gray).
- Z-order set to -10.

Plot three lines with distinct characteristics:

- The first line (Blue signal) should use color 'C0', linewidth of 2.5, and be placed at z-order 10.
- The second line (Orange signal) should use color 'C1' and linewidth of 2.5.
- The third line should consist of scatter markers at every third point, with no linewidth, markersize of 9, marker style 's' (square), marker face color 'none', marker edge color 'C4', and marker edge width of 2.5.

Set the title "Anatomy of a figure" and axis labels "x Axis label" and "y Axis label" with specific font sizes:

- Title font size should be 20.
- Axis label font sizes should be 14.

Add a legend with these specifications:

- Positioned at the "upper right".
- Font size set to 14.

Annotate the figure, including tick labels, axes labels, grid, etc., using circles, text, and code snippets at specified coordinates. The circles should have a radius of 0.15, a border color defined by the royal_blue variable with an alpha of 0.6, and a white, non-filled center. Text annotations should be in both royal blue and black colors, with bold and italic styles.

CHARTUIE-8K Example

Create a bar chart showing the GDP compositions of agriculture, industry, and services for the top 5 countries by GDP. Highlight the country with the highest percentage of each sector. Add annotations to display exact values for each bar segment, focusing particularly on the leading country in each sector.

D Human Study Details

D.1 Compensation and Qualification

We use paid crowdsource surveyors to conduct our human study. We use the same platform as [Gudibande et al. \(2024\)](#). To ensure quality responses, all voluntary participants are paid \$14 per hour, well above the U.S. federal minimum wage of \$7.25 ([Henderson, 2024](#)). We ensure that surveyors feel safe to express freely by clearly indicating: "feel free to be wholly honest". Moreover, we enforce five requirements to participate in our study: (i) Masters², (ii) > 80% Approval Rate, (iii) > 50 Approved, (iv) Casual experience using LLMs, (v) Survey with a large screen (≥ 10 inches, ≥ 25.4 cm). For privacy, we do not gather any other unnecessary data.

D.2 Sanity Check

Following [Yang et al. \(2023\)](#) and [Xu et al. \(2024a\)](#), we include multiple sanity check questions within the survey for reliability. These sanity checks are solidified *a priori*, and are *never* changed.

CHARTUIE-8K. In CHARTUIE-8K, we manually remove LLM initial query that clearly suggests that the surveyor did not understand the survey instructions or put in virtually zero time. For transparency, we report *all* the sanity check fail cases on our open-source [github repository](#).

CHARTAF. For the CHARTAF human study, two sanity checks filter very poor quality surveyor responses. First, we remove surveys that were completed in 50% of the lower-bound completion time. We time the duration it takes to complete a survey, and create a lower- and upper-bound to signal to potential surveys how long our study takes. For example, 10 to 16 minutes. In this case, surveys that were completed within 5 minutes would be removed as this would indicate that *very low effort* was put into the survey. Second, a sanity check question was included, presenting an image option intentionally unrelated to the user instructions. Survey responses that prefer (or strongly prefer) this option is removed as it indicates that the surveyor does *not* understand the task or is putting in very little effort.

D.3 CHARTAF Human Study

We provide the precise instructions provided to the surveyors below. We further provide exact screenshots in Fig. 7 and 8.

CHARTAF Human Study (Part 1)

User Instruction
USER REQUEST
[insert initial instruction]

QA WITH USER
[insert QA]

Referring to the User Instruction, carefully select the **better** chart.
[insert generated image by f]
[insert generated image by
post-CHARTAF-feedback
or post-ChartX-feedback
or post-ChartX+A-CoT+SC-feedback
or Sanity Check
or post-ChatEval-feedback
or post-NF-feedback
or post-MatPlotBench-feedback
or post-MatPlotBench+A-CoT+SC-feedback
or post-Plot2Code-feedback

²Masters are "Workers who have demonstrated excellence across a wide range of tasks are awarded the Masters Qualification. Masters must continue to pass our statistical monitoring to retain the Masters Qualification."

or post-Plot2Code+A-CoT+SC-feedback
or post-Skip+A-CoT-feedback]
[insert preference score options]

CHARTAF Human Study (Part 2)

User Instruction

USER REQUEST

[insert initial instruction]

QA WITH USER

[insert QA]

Referring to the **User Instruction**, carefully select the better chart.

[insert generated image by pre-TTS]

[insert generated image by post-TTS]

[insert preference score options]

D.4 CHARTUIE-8K Human Study

We provide the precise instructions provided to the surveyors below. We further provide screenshots in Fig. 9 and 10.

CHARTUIE-8K Human Study (Part 1)

- Looking at the following sample image, imagine you want to ask Large Language Model (like ChatGPT) to draw a chart like this image:

[insert example chart image]

- Please write how you would prompt (instruct) an LLM
- (You can assume that you have already provided the dataset to the LLM.)

[insert text form field]

CHARTUIE-8K Human Study (Part 2)

- After the user has given an initial set of instructions, the LLM may be inclined to ask some questions to make sure it clarifies on any lack of detail or subjective instructions.
- The following is an example of the LLM's response to the user's initial instructions.

[insert sample initial instruction and questions]

- If you were the user, would you prefer the LLM respond to you with, for example, the **clarification questions** above?

[insert binary option (Yes or No)]



- After the user has given an initial set of instructions, the LLM may be inclined to ask some questions to make sure it clarifies on any lack of detail or subjective instructions.
- The following is an example of the LLM's response to the user's initial instructions.

[insert sample initial instruction, questions, and answers for the questions]

- Do you think that the **user response** above is **realistic**?

[insert binary option (Yes or No)]

* 8. Referring to the **User Instruction**, carefully select the better chart



Size of each tile represents the absolute value of the measurement.
Colors represent different body parts.

Strongly prefer Top Prefer Top Indifferent Prefer Bottom Strongly prefer Bottom

Which chart is better?

Figure 7: ChartAF Human Study (Part 1) Example

User Instruction

USER REQUEST

Create a radar chart to compare different factors affecting students' adaptability in online education. Include variables like 'Gender', 'Age', 'Education Level', 'Institution Type', 'IT Student', 'Location', 'Financial Condition', 'Internet Type', etc.

QA WITH USER


Q: Should the radar chart include a legend to identify each variable?

A: Yes, please include a legend.

Q: Would you prefer shaded areas or lines to represent the different factors on the radar chart?

A: I would prefer shaded areas.

* 12. Choose better Chart considering **User Instruction** carefully.




Strongly prefer Bottom Prefer Bottom Indifferent Prefer Top Strongly prefer Top

Which chart is better?

Figure 8: ChartAF Human Study (Part 2) Example

* 1. Looking at the following sample image, imagine you want to ask Large Language Model (like ChatGPT) to draw a chart like this image:



Please write how you would prompt (instruct) an LLM

(You can assume that you have already provided the data set to the LLM.)

Figure 9: ChartUIE-8K Human Study (Part 1) Example

* 2.

- After the user has given an initial set of instructions, the LLM may be inclined to ask some questions to make sure it clarifies on any lack of detail or subjective instructions.
- The following is an example of the LLM's response to the user's initial instructions.

Please create a histogram for the Crop Recommendation Dataset. Specifically, plot histograms for the distributions of Nitrogen (N), Phosphorus (P), and Potassium (K) levels. Ensure each histogram is clearly labeled and shows the frequency of each nutrient level to help visualizing nutrient distribution in various crops.



1. Should the histograms have grid lines for better readability?
2. Would you like to include a title for each histogram? If so, please specify.
3. Do you prefer the histograms to be displayed in a single row, column, or in a grid layout?
4. Are there any specific axis ranges you would like for the histograms? If so, please specify.

- If you were the user, would you prefer the LLM respond to you with, for example, the **clarification questions** above?

Yes

No

* 3.

- After the user has given an initial set of instructions, the LLM may be inclined to ask some questions to make sure it clarifies on any lack of detail or subjective instructions.
- The following is an example of the LLM's response to the user's initial instructions.

Please create a histogram for the Crop Recommendation Dataset. Specifically, plot histograms for the distributions of Nitrogen (N), Phosphorus (P), and Potassium (K) levels. Ensure each histogram is clearly labeled and shows the frequency of each nutrient level to help visualizing nutrient distribution in various crops.



1. Should the histograms have grid lines for better readability?
2. Would you like to include a title for each histogram? If so, please specify.
3. Do you prefer the histograms to be displayed in a single row, column, or in a grid layout?
4. Are there any specific axis ranges you would like for the histograms? If so, please specify.

1. Yes, the histograms should have grid lines for better readability.
2. Yes, please include a title for each histogram: "Distribution of Nitrogen (N) Levels", "Distribution of Phosphorus (P) Levels", and "Distribution of Potassium (K) Levels".
3. A grid layout would be preferred.
4. I do not have any specific axis ranges in mind.



- Do you think that the **user response** above is **realistic**?

Yes

No

Figure 10: ChartUIE-8K Human Study (Part 2) Example

E LLM Configurations

Table 5 shows the LLM configurations we used. These configurations are determined *a priori* to all experiments, and never changed.

	Model	Temp.	Max Tokens
GPT 4o	gpt-4o-2024-08-06	1	4096
Claude 3.5 Sonnet	claude-3-5-sonnet-20240620	1	4096
Llama 3.1 70B	llama-3.1-70b-versatile	1	4096
Gemma 2 27B	gemma-2-27b-it	1	2048

Table 5: LLM Configurations

F Code Error Rate

Fig. 11 shows the mean execution error rates of the four tested LLMs. This is an inherent limitation of the original chart-generating LLM and does not relate to our method.




Figure 11: Execution error rates of the tested LLMs

G Icon Attribution

All icons are from flaticon.com except Retain(cycle-100) is from icons8.com. The icons and their respective authors are in Table 6. The license to these icons belongs to their respective owners. We use them for non-profit academic use, with attribution as requested by the distributors.

Category (Icon)	Author(s)
Retain (cycle-100)	Icons8
Discard (remove)	Pixel perfect
Add (add)	Arkinasi
Task (checklist)	juicy_fish
Purpose (target)	Freepik
Audience (users-avatar)	SeyDesigner
Basic Criteria (list)	Smashicons
Addition (add)	Digby Garrett
Query (question)	Md Tanvirul Haque
Binarize (file-management)	iconsmind
Social Science (network)	Freepik
Biology (book)	Freepik
Box Chart (chart)	Freepik
Stacked Chart (column-chart)	Freepik
LLM (big-data)	LAFS
AI (ai)	Freepik
Number 1 (circle-1)	iconographics
Number 2 (circle-2)	iconographics
Number 3 (number-3)	BS Editing
LLM to Evaluate (dashboard)	xnimrodx
Data Set Filter (database-structure)	Uniconlabs
Chart Annotations (notes)	Freepik
Data Set Crawl (bug)	Freepik
Coding (code-review)	Freepik
Evaluate (criteria)	cah nggunug
Edit (modified)	Uniconlabs

Table 6: Icons and Respective Authors