

Visualization Literacy of Multimodal Large Language Models: A Comparative Study

Zhimin Li*
University of Utah

Haichao Miao†
Lawrence Livermore
National Laboratory

Valerio Pascucci ‡
University of Utah

Shusen Liu §
Lawrence Livermore
National Laboratory

ABSTRACT

The recent introduction of multimodal large language models (MLLMs) combine the inherent power of large language models (LLMs) with the renewed capabilities to reason about the multimodal context. The potential usage scenarios for MLLMs significantly outpace their text-only counterparts. Many recent works in visualization have demonstrated MLLMs’ capability to understand and interpret visualization results and explain the content of the visualization to users in natural language. In the machine learning community, the general vision capabilities of MLLMs have been evaluated and tested through various visual understanding benchmarks. However, the ability of MLLMs to accomplish specific visualization tasks based on visual perception has not been properly explored and evaluated, particularly, from a visualization-centric perspective.

In this work, we aim to fill the gap by utilizing the concept of visualization literacy to evaluate MLLMs. We assess MLLMs’ performance over two established visualization literacy evaluation datasets (VLAT and mini-VLAT). Under the framework of visualization literacy, we develop a general setup to compare different multimodal large language models (e.g., GPT4-o, Claude 3 Opus, Gemini 1.5 Pro) as well as against existing human baselines. Our study demonstrates MLLMs’ competitive performance in visualization literacy, where they outperform humans in certain tasks such as identifying correlations, clusters, and hierarchical structures.

Index Terms: Multimodal Large Language Model, Visualization Literacy

1 INTRODUCTION

Recently, due to the advance of generative neural network models such as large language models (LLMs), we can easily produce visualization based on natural language instruction by generating visualization code. However, due to the modality of model input, LLMs can not easily understand the generated visualization in image form. That limitation changes with the introduction of multi-modal capability, i.e., multi-modal large language models (MLLMs), in which the LLMs also obtain the ability to understand vision input, including visualization results [35]. Such capabilities facilitate novel usage of MLLMs, e.g., developing agents that can accomplish potentially complex visualization tasks [19].

However, all these potential use cases hinged on the model’s ability to understand the visualization, which is currently not well studied, particularly, regarding the extent of their abilities and limitations. As a result, a systematic evaluation of the ability of MLLMs to understand the visualization is the fundamental step we need

before applying them in downstream applications. Several benchmarks have been introduced in the machine learning community for evaluating the model’s capabilities for chart and infographic understanding [20, 21, 14]. They provide valuable assessment regarding the general capability for understanding information in graphics form, however, due to their unstructured nature, it is still hard to understand the model’s performance on specific types of visual encoding or visualization tasks. Moreover, these datasets are not designed with visualization evaluation in mind, therefore the example selection process is likely non-discriminative in terms of poor visualization design or includes many illustrations instead of proper visualization.

In this work, we aim to fill this gap by obtaining a fine-grain assessment of MLLM’s behavior on a smaller yet well-curated dataset that covers major visual encoding and visualization tasks. We achieve this by bringing the concept of visualization literacy and the associated tests for evaluating the MLLM’s capability for understanding visualization. We compared multiple state-of-the-art models, as well as against human performance from previous visualization literacy research on the same test.

Literacy often refers to the ability to read and write. In a broader sense, it encompasses the capacity to effectively understand, interpret, and communicate non-verbal information in written form. Following a similar definition, visualization researchers brought the concept of visualization literacy, which describes a person’s ability to interpret visual communication through visualizations such as various charts and graphs. A curated set of questions has been proposed to cover different types of visual encoding with varying types of questions. Despite the small size of these questions and examples, they aim for a broad coverage both in terms of visual encoding and the type of queries. Most notably, VLAT and Mini-VLAT, present well-recognized and well-established datasets [24, 25, 16] for visualization literacy assessment.

This work focuses on investigating how well MLLM models can perform on these visualization literacy tests, how their behaviors are compared with humans, and what are their strengths and weaknesses. Through our study, we find that MLLMs are effective at visualization tasks such as correlation trends, cluster analysis, and identifying hierarchical structures. Despite the competitive performance of MLLMs over visualization literacy, we observe that humans and MLLMs demonstrate distinct failure patterns. Our key contributions are summarized below:

- Introduce the concept of visualization literacy for evaluating MLLMs’ capabilities for understanding visualization.
- Present a comparative evaluation of multiple state-of-art MLLMs over the Mini-VLAT and VLAT visualization literacy test, as well as their performance against an open-source human evaluation result on the same test.
- Provide the accompanying error analysis for uncovering the limitation of current MLLMs for visualization tasks and encodings.

*e-mail: zhimin@sci.utah.edu

†e-mail: miao1@llnl.gov

‡e-mail: pascucci@sci.utah.edu

§e-mail: liu42@llnl.gov

2 RELATED WORK

2.1 Visualization competency test

A human’s ability to understand visualization and visual encoding can vary greatly depending on their education level and prior exposure. Börner et al. [4] evaluate the 273 museum visitors with different backgrounds including youth and adults. Their evaluation involves 20 visualizations that show up in newspapers and textbooks. They find that many interview candidates cannot understand and interpret the presented visualization. The observation implies information visualization is not well-accessed by the population. Various visualization competency tests have been introduced as part of visualization education efforts. Alper et al. [1] explored how visualization is taught at elementary school and interviewed the school teacher about the challenge of visualization teaching. A web application is developed at the end to improve the visualization literacy of children. E. Firat et al. [9] develop the treemap literacy which specifically focuses on evaluating the human’s performance over treemap visualization. VLAT [16] is a well-established test used to measure visualization literacy. The overall evaluation contains 12 visualizations and 53 questions. Pandey and Ottley [25] shorted the test into a Mini-VLAT dataset which only contains 12 questions. Carolina et al. [24] replicate the evaluation study of the VLAT evaluation and give a more detailed understanding of the barriers during the visualization interpretation process. The discussion of the previous visualization competency research mainly assesses humans’ capabilities. However, in our work, the main focus is MLLM. We compare the performance of MLLMs and discuss the reason for their mistakes.

2.2 MLLM and LLM for Visualization

The MLLMs and LLMs have demonstrated outstanding abilities to understand and solve a wide range of practical problems [35]. Many research efforts have utilized the generative capability of large language model [29, 32, 17] to create visualization through code generation or to interpret and manipulate the visualization with the SVG format. Chen et al. [7] evaluate the performance of GPT3.5 and GPT4 over a data visualization course and find out that the LLM can score 80% of the assignment. Such a finding motivates the new requirement of visualization education design in college. In the application of Rich Screen Reader, Zong et al. [39] use LLM to describe visualization to people with impaired vision. Yang et al. [34] present a review study of how the foundational model will help visualization over the previous literature. Liu. et al. [19] proposed the concept of the autonomous visualization agent, which utilizes the visual perception ability of MLLMs to directly understand user intention adjust visualization parameters, and help solve visualization tasks autonomously. Compared with our study, most of these works do not involve the new emerging perception ability of MLLMs, and they often only focus on models from a single source (e.g., OpenAI). Instead, our study focuses on the vision aspect of MLLMs and evaluates their visualization literacy.

2.3 Machine Perception for Visualization

Visual perception is a fundamental piece of the visualization system design. Due to the complexity of human evaluation, many researchers aim to use machine learning models (e.g., neural networks) to assist or even replace human evaluation. Convolution neural network is a potential candidate for such a task given that its initial design is inspired by the mechanism of humans’ visual cortex [15]. Giovannangeli, Loann, et al [10] use this idea to predict the performance of node-link diagrams and adjacency matrix diagrams over humans with a neural network evaluation. Haehn et al. [11] study the analysis ability of convolution neural networks on visualization tasks by evaluating their performance over elementary perceptual tasks from Cleveland and McGill [8]. Their study concludes that CNN is not an ideal model for human perception modeling

which is similar to previous discovering [5]. Yan et al. [33] study over twenty-nine neural network architectures to predict human behavior over the correlation comparison in scatter plots. There are more works [10, 38] under this path to study the correlation between neural networks and human performance. Compared with previous work, our study focuses on more powerful Multimodal large language models that are pre-trained with significantly larger datasets than the traditional network models (e.g., CNN pre-trained with ImageNet data). Moreover, our evaluation focuses on understanding MLLMs’ ability in visualization literacy in general instead of a specific visualization or task.

3 BACKGROUND

3.1 Multimodal Vision Language Models

Recently, models that accept multiple modalities as input are gaining increasing popularity. From the early work, e.g., CLIP [26], that jointly embedded image and text to models designed for input in one modality and output in another, such as text-to-image [28, 27] and image-to-text [37] models. In the context of this work, we focus on MLLM [36] that can understand both image and text inputs for text generation. LLMs are often referred to as foundation models due to their extensive capabilities and scale. They often possess a wide range of knowledge, allowing them to understand implicit context and common sense. Since humans interact with their environment through multi-modal sensory input, the evolution from text-only LLM systems to multi-modal LLMs significantly increases the potential usage cases for LLMs, prompting many applications [35]. Since the initial introduction of GPT4-V, many state-of-the-art MLLMs have been introduced such as the recent update to GPT4-o from OpenAI, multimodal support for Claude 3 [3] from Anthropic, and Gemini family from Google. On the open weight model side, many attempts have been made to develop similar capabilities, such as LLAVA [18], and CogVLM [30].

3.2 Visualization Literacy

In the Merriam-Webster dictionary, literacy is defined as “the ability to read and write” Visualization literacy is often described as the ability to understand and extract information from data visualization. Understanding visualization literacy has attracted a broad attention [16, 6, 1, 4]. In particular, the VLAT data is a popular and well-established dataset that is used to measure visualization literacy [25, 24] of humans. The dataset contains 12 different visualizations with 53 questions. In this study, we collect our visualization literacy test from the original authors’ website¹. Mini-VLAT is a short version of the VLAT dataset which has the same 12 visualizations but only presents 12 questions. We collect this dataset from the original authors’ repository². Previous research [25] shows that the behavior of humans in the Mini-VLAT dataset shows similar behavior patterns over the VLAT dataset.

4 EXPERIMENT DESIGN

To perform a meaningful comparison, we choose the best MLLMs from the frontier labs (OpenAI, Google, and Anthropic), namely, GPT4-o, Gemini Pro, and Claude 3 Opus, as the candidate models and evaluate their performance. They are currently the most powerful models from the most dominant companies in the LLM space and often demonstrate state-of-the-art performance over different evaluation benchmarks in the current literature.

The experiment tests each multimodal language model on the Mini-VLAT and VLAT datasets. It is important to notice that LLM can give different answers to the same question in separate runs (due to the random seed). Therefore, each LLM will be tested 10 times on the same question separately. The final result for each

¹<https://www.bckwon.com/publication/vlat/>

²<https://github.com/washuvis/minivlat/blob/main/MiniVlatQUESTIONS.pdf>




Figure 1: The result shows the performance of different multimodal large language models over the Mini-VLAT dataset.

question is reported as the average score of all 10 runs (each correct answer gets a score of 1.0, whereas the incorrect answer gets a score of 0.0). The overall testing and scoring process is fully automatic and the model query is through Python script which is implemented with the API of OpenAI, Google, and Anthropic. For LLM such as GPT4-o, it costs 0.26\$ and less than 5 minutes to answer all 53 questions from the VLAT dataset. As a reference, hiring a human to perform the same task [24] takes around a half hour to finish the test.

How to ask a question can affect the quality of its answer. Similarly, how to ask an LLM question, which is often considered as the prompt, can significantly affect the performance of an LLM [31]. In our setup, we aim to use a basic but functional prompt that introduces as little interference or perturbation to the original model’s capability as possible. Optimizing prompts can be an extensive undertaking and can also lead to implicit bias toward one model or another. As a result, in our experiment, we did not optimize the prompt or employ enhanced reasoning approaches such as chain-of-thoughts [31]. We use the same prompt for all models with the following format. An example prompt for one question given a visualization:

What was the price of a barrel of oil in February 2015? Answer: 1) \$57.36 2) \$47.82 3) \$50.24 4) \$39.72. Please answer with Answer: 1) - 4) and Why:

Beyond the question and answer, the response format helps us produce more consistent answers for automatic scoring. The “why” is included to provide a bit more context for the answer for our subsequent analysis.

We use the open-source human evaluation result³ for comparison. This dataset is collected by the author from the Mini-VLAT to replicate the previous VLAT study. The total number of people who finish the test is 199 participants. Most of them have high school diplomas. 19% have a Master’s degree and around 40% have a bachelor’s degree. During the question and answer process, each question needs to be completed in 25 seconds. The author claims that the overall score of this replicate study is similar to the original VLAT study [16]. More detail about the study can be found in the original experiment [25].

5 EVALUATION

5.1 Evaluation on Mini-VLAT

We perform a pilot evaluation study with the Mini-VLAT dataset to compare three state-of-the-art MLLMs’ performance alongside a random selection baseline to help gauge the models’ performance

³<https://osf.io/dy67k/>

MLLM	GPT4-o	Claude-opus	Gemini-1.5-pro	RB
best	0.833	0.75	0.666	–
worst	0.583	0.333	0.333	–
mean	0.741	0.483	0.508	0.35

Table 1: The performance of different MLLMs over the Mini-VLAT datasets. The best/worst performance measures the highest/lowest single run score out of all runs. The average score is the mean of all runs.

(the questions are with different numbers of choices, so the score of random guesses is not obvious). The summary of the result is described in TABLE 1. The best and average performance of each model is significantly better than the random baseline. In this case, the GPT4-o, which was recently released by OpenAI, has the best statistic over all three measurements. Gemini has the second performance on average, while the Claude model has the advantage in terms of its best run. The detailed performance breakdown of each question is described in Fig. 1. All models correctly answer questions for the treemap and scatter plot visualization. They also perform well over the Choropleth chart [2], Line chart, and Bar chart. The GPT4-o model performs better than others in the Histogram chart and Area chart.

5.2 Error Analysis

To better understand the model’s behavior, it is crucial to investigate beyond just summary statistics and understand where and how the model fails. Existing literature has discussed the failure reasons from the human visualization literacy tests [24]. In this section, we will dive into the details and try to understand why an MLLM fails a question. Our discussion focuses on the question in the Mini-VLAT dataset.

Given the overall performance of GPT4-o outperforms other models, we mainly discuss the error analysis of the GPT4-o model. Specifically, our analysis focuses on the 100% Stacked Bar chart, Pie Chart, and Stacked Bar Chart. The success rate of GPT4-o over the Stack Area Chart is also low. However, this question is related to the value-retrieving task which is the same as the Pie Chart and Stacked Bar Chart. So, we will not discuss the details here. During the discussion, we hypothesize the reason for failure and try to verify the hypothesis by adding annotation in an attempt to remove potential confusion and improve the success rate. This process will iterate 2 to 3 times until the MLLM can pass all the question tests.

5.2.1 Confusion from Color-Semantic Association

The color chosen to encode the visualization may affect the decision of the LLM model. Fig. 2 displays a 100% stack bar chart, and the question is:

Ask: “Question: Which country has the lowest proportion of Gold medals? Answers: 1) Great Britain, 2) USA, 3) Japan, 4) Australia. Please answer with Answer: 1) - 4) and Why:”

Response: “Answer: 3) Japan Why: The chart shows that Japan has the lowest proportion of Gold medals compared to the other countries, as indicated by the smallest relative height of the gold segment in the stacked bar for Japan.”

For this question, Japan and Australia are frequent answers and each occurs rate is 4 out of 10 run, i.e., 4/10. The rest of the answers are 1/10 U.S.A and 1/10 Great Britain. Here, Great Britain is the correct option. However, Japan has the highest percentage of golden metal, and the answer from MLLM does not seem to be reasonable. In the visualization (fig. 2 (1)), the brownish golden color and the orange color not only look similar, but the color indicates the brown metal is more likely to be referred to as golden color compared to the orange that represents the golden metal in the plot. From the explanation of the answer, we can tell that the MLLM model considers the brownish golden color as the golden

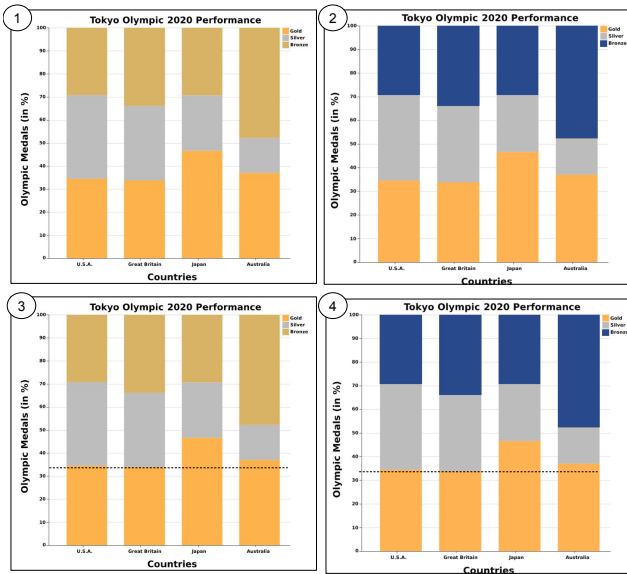


Figure 2: The language model may be confused by the ambiguous colors and resort to their prior knowledge root in color-semantic association [23], which leads to incorrect answers.

percentage and considers Japan as the final answer. This explanation shows that MLLM is confused about what color should represent the golden metals, and potentially the existing color-semantic association may lead to such a mistake, a similar phenomenon has been observed in MLLMs [23].

To verify the hypothesis that MLLM may be confused with golden color, we updated the visualization in two different ways. In Fig. 2 (3), we add a dashed line to help the MLLMs to compare the golden metals over different countries. With the assistance of the dashed line, the response score of the MLLM is 5/10 Japan and 5/10 Great Britain. The success rate of the model is improved but the MLLM is still confused with the color, and Japan is still a frequent answer in the updated visualization. In Fig. 2 (2), we updated the color to dark blue and performed the same experiments 10 times to measure the overall success rate. The final statistic is 7/10 Australia and 3/10 Great Britain. The final answer is significantly shifted with the new visualization. In the visualization, the differences between Great Britain and Australia are still minor which is a potential reason why the MLLM failed the task in the new visualization. In the end, we apply both updates in the visualization in Fig. 2 (4) and experiment again. The response is 10/10 Great Britain. This experiment helps reveal the potential reason that an MLLM fails a question here is the color and the minor difference between the two categories during comparison.

5.2.2 Poor Performance at Value Retrieval

Furthermore, we investigate the relatively poor performance for value retrieval tasks, here we take the Stacked Bar Chart as an example. The task is to measure the length of a rectangle. In Fig. 3 (1), the distribution of the answer out of 10 runs is as follows, 5/10 the answer is 7.5, 3/10 is 6.1, and 2/10 is 5.2.

In Fig. 3 (2), during the exploration and error analysis, we worried the lack of contrasts between colors may cause such a mistake during the decision process. We updated the visualization color of the peanut to make the overall comparison more distinguishable, but the final answer shows that 9/10 is 7.5, and 1/10 is 6.1. The answers from the MLLM become more stable but the performance decreases. So likely, the color is not the main contributor to the mistakes. Now, let's take a closer look at the question and response.

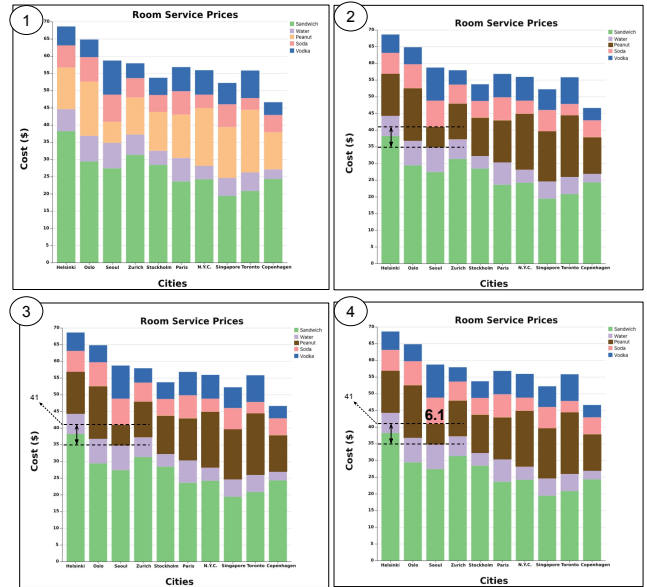


Figure 3: The language model has limited ability to retrieve value.

Ask: "Question: What is the cost of peanuts in Seoul? Answers: 1) 5.2, 2) 6.1, 3) 7.5, 4) 4.5. Please answer with Answer: 1) - 4) and Why:"

Response: "Answer: 3) 7.5 Why: To find the cost of peanuts in Seoul, we need to look at the height of the brown section of the stacked bar for Seoul. The cost for each item can be deduced by noting the cumulative height of the sections up to the top of the brown segment. The green section (Sandwich) ends at about 30. The purple section (Water) ends at about 35. The brown section (Peanut) ends at approximately 42.5. The height of the brown section (Peanut) can be calculated as the difference between the cumulative heights of the brown and purple sections: $42.5 - 35 = 7.5$. Therefore, the cost of peanuts in Seoul is 7.5."

From the response of GPT4-o, we can tell that the MLLMs consider the top rectangle to be aligned with 42.5. However, the axis ticks tell us that the top rectangle has not reached the half region between 40 and 45, and the aligned value should be smaller than 42.5. The GPT4-o fails to retrieve a correct value to do the analysis. To help GPT4-o better understand this visualization, in Fig. 3 (3), we add additional annotations (labeling the correct axis value) in the visualization to help MLLMs understand the visualization and retrieve the value. With the assistance of the new visualization, the response of the MLLMs changed to 7.5 (5/10 runs), and 6.1 (5/10 runs). We consider the question in Fig. 3 (3) to be simplified but the model still demonstrates a limited ability to answer this question correctly. In the end, we have to do the math and add the measurement result to the visualization (Fig. 3 (4)) to help the MLLM pass all the tests.

5.2.3 Difficulty in Interpreting Pie-Chart Visual Encoding

Previous cases demonstrate the potential limitation of the MLLMs over the Stack Bar Chart. The last error analysis is with the Pie chart. With the visualization in Fig. 4 (1), the model responds 17.6% (4/10) which is the correct answer. However, the rest of the answer is 25.3% and the detailed decision process of the model can be found in the following response.

Ask: "Question: What is the approximate global smartphone market share of Samsung? Answers: 1) 17.6%, 2) 25.3%, 3) 10.9%, 4) 35.2% Please answer with Answer: 1) - 4) and Why:"

Response: "Answer: 2) Why: By observing the pie chart, the




Figure 4: The language model has limited ability to retrieve the arc visual encoding value.

approximate market share of Samsung is a bit more than a quarter of the total chart. Therefore, it aligns most closely with the option 25.3%.

The first hypothesis is that the model missing a value reference to understand the percentage of the area. To verify this, we add a dashed line in the visualization to clear-cut the pie chart into 4 pies, and each represents 25%. The new visualization is updated in Fig. 4 (2). However, the results show only 4/10 runs give correct answers. We added additional references in the visualization to identify the region belonging to Samsung, but the final response in Fig. 4 (3) still only shows correct answers in 4/10 runs. In the end, we have to label this region with 17% to achieve a 10/10 success rate.

5.3 Evaluation on VLAT Dataset

Given the mostly favorable performance of MLLM over the Mini-VLAT dataset (particularly for GPT4-o), we perform the same evaluation over the VLAT dataset with 53 questions to measure the performance of the models on a wider set of tasks and visual encodings. During the evaluation, one interesting question to explore is whether there is a clear distinguishing pattern between human visualization literacy and MLLM visualization literacy, e.g., where MLLMs align with human behaviors, and where they diverge.

Table 2 provides the summary of the evaluation result over MLLMs and humans over nine tasks. The best performance in each task is highlighted in bold text. RB represents the random baseline, where we random select an option in each question. Both humans and MLLMs perform better than the random baseline. Comparing the average performance of MLLMs and humans, humans perform better than the MLLMs on average. The maximum performance of a human is much better than that of a MLLM. For the minimum performance, MLLM appears to be better than a human. However, the reported minimal human performance is close to random selection. Since the answers are crowd-sourced the minimum performance estimation may represent a lack of effort from the participant rather than an accurate estimation of human capabilities.

From the comparison, we observe that humans outperform MLLMs on **Characterize-Distribution**, **Retrieve-Value**, and

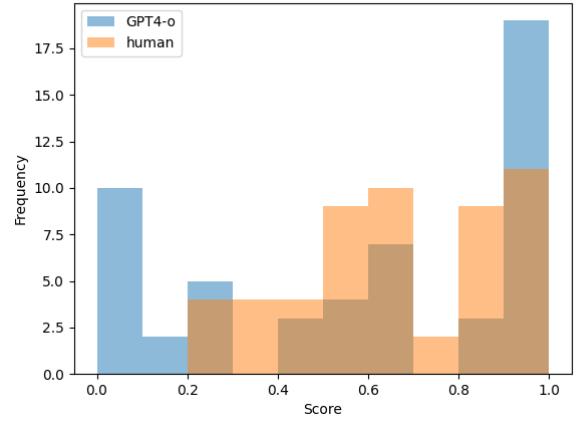


Figure 5: Average score distribution of 53 questions from human and GPT4-o.

Find-Extremum with a significant margin. For the rest of the tasks, MLLMs outperform humans. Especially, MLLMs perform better than humans on tasks such as **Find Correlations/Trends**, **Find Clusters**, and **Identify the Hierarchical Structure**. If we consider the visualization task as a basic unit to compare the behavior of humans and MLLMs, we can tell from the correlation matrix in Fig. 6 that humans' response to visualization has a strong correlation with the MLLMs such as GPT4-o and Claude. Task-based correlation between humans and neural network models is also discovered in previous literature [10] over graph visualization task. However, this statement needs more studies to be verified in future studies. In the meantime, MLLMs also have a strong correlation with each other except Gemini which demonstrates a minor correlation.

Table 3 displays the detailed performance of each question, and its corresponding visual encoding and task type. With 14 over 53 questions, humans outperform MLLMs. However, humans have the best average performance in 6 out of 12 visual encoding types. In the rest of the cases, the MLLM model outperforms humans. From the evaluation result, we identify some interesting observations.

Humans have a more stable performance than MLLMs. The value range of humans' average performance ranges from 0.2 to 1.0, and the overall distribution is coherent. During the evaluation process, an MLLM may display a 0% success rate when answering certain questions, such phenomena do not occur for humans. On the other hand, MLLMs may show excellent performance in certain tasks such as Line Chart item 2, item 3, and item 4. All three models have a 100% success rate to answer these questions. From the score distribution of humans and MLLMs in Fig. 5, we can conclude that the performance of humans is more stable than GPT4-o. A similar observation was also discovered from previous work [7] with a text-based GPT model. However, MLLMs have more extreme cases that either perform the test well or badly. Fig. 5 only compares GPT4-o with humans but the other two models also show similar properties. This observation led to another question whether there is a correlation between humans and MLLMs over per-question scores. The answer can be found in Fig. 7 that the correlation between humans and MLLMs is not strong. The correlation between MLLMs is also weak.

From the evaluation result of the Mini-VLAT and VLAT datasets, the performance of MLLMs demonstrates a competitive ability in visualization literacy when compared with humans. We can tell that MLLMs may have certain similarities with humans'

Task/LLM	GPT4-o	Claude-opus	Gemini-1.5-pro	human	RB
Retrieve Value(13)	0.461	0.338	0.361	0.621	-
Find Extremum (12)	0.683	0.491	0.675	0.782	-
Determine Range (5)	0.5	0.48	0.559	0.492	-
Find Correlations/Trends (5)	0.82	0.820	0.840	0.774	-
Make Comparisons (12)	0.575	0.75	0.525	0.648	-
Characterize Distribution (1)	0.0	0.0	0.266	0.457	-
Find Anomalies (2)	0.35	0.2	0.2	0.485	-
Find Clusters (2)	0.85	0.85	0.05	0.751	-
identify the hierarchical structure (1)	1.0	1.0	0.9	0.879	-
All (53) maximum	0.660	0.622	0.566	0.943	-
All (53) minimum	0.547	0.452	0.453	0.339	-
All (53) mean	0.586	0.545	0.518	0.667	0.33

Table 2: The performance of four MLLMs and humans on the VLAT datasets.

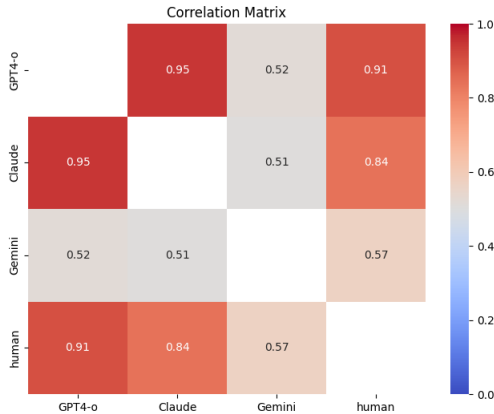


Figure 6: Task-based performance correlation analysis between MLLMs and humans

behavior over visualization literacy tests but such similarities may not be particularly strong.

6 DISCUSSION AND CONCLUSION

The new emerging visual perception ability of the MLLMs brings new opportunities and challenges to the visualization community. How this new ability will change the way humans interact with visualization is an exciting research direction in the expected future. A critical step to utilize this new technique is to rely on the proper evaluation of the MLLMs to reveal their advantages and limitations. In this study, we leverage the Mini-VLAT and VLAT datasets, the popular dataset for visualization literacy measurement, to understand MLLMs’ capability to accomplish basic visualization understanding tasks. Our evaluation tells us that MLLMs already demonstrate competitive performance over these visualization literacy tests, and can perform certain tasks such as correlation/trend analysis, and treemap analysis very well. MLLMs and humans demonstrate different behavior patterns in their question score distribution and the question-based correlation may not be particularly strong. Our error analysis with the MLLMs also reveals that these models still have quite a few limitations for the reliable interpretation of visualization and visual encodings. The MLLMs still struggle to interpret the basic elements such as the portion of a pie chart with additional annotation in the visualization.

Our study mainly focuses on the basic prompt to understand the performance of the multi-model large language model. Prompt engineering can play a significant role in the performance of the model. It would be interesting to explore how different prompting




Figure 7: Question-based performance correlation analysis between MLLMs and humans

techniques [31] for MLLMs can affect the overall performance, and whether they are effective for vision tasks. We believe our study represents one of the necessary steps for revealing the capability of MLLM for visualization interpretation. There is a broad spectrum of research that needs to be done to evaluate the ability of MLLMs such as how MLLMs react to different visual channels [22, 12], and how MLLMs react to more advanced visualization techniques such as parallel coordinate [13] visualization.

ACKNOWLEDGMENTS

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. This work is partially supported by LDRD 23-ERD-029 and 23-SI-003. The manuscript is reviewed and released under LLNL-PROC-865825.

Item ID	Visualization	Task	GPT4-o	Claude-opus	Gemini-1.5-pro	Human
item1	Line Chart (1)	Retrieve Value	0.2	0.0	0.0	0.693
item2		Find Extremum	1.0	1.0	1.0	0.964
item3		Determine Range	1.0	1.0	1.0	0.542
item4		Find Correlations/Trends	1.0	1.0	1.0	0.954
item5		Make Comparisons	0.0	0.5	0.2	0.613
mean				0.64	0.7	0.64
item1	Bar Chart(2)	Retrieve Value	0.4	0.3	0.1	0.839
item2		Find Extremum	1.0	0.8	1.0	0.959
item3		Determine Range	0.0	0.3	0.0	0.593
item4		Make Comparisons	0.4	0.5	0.8	0.261
mean				0.449	0.475	0.475
item1	Stacked Bar Chart(3)	Retrieve Value	0.2	0.0	0.0	0.884
item2		Retrieve Value	0.5	0.6	0.1	0.402
item3		Find Extremum	0.0	0.6	0.0	0.703
item5		Make Comparisons	0.0	1.0	0.0	0.568
item6		Make Comparisons	0.4	0.7	0.9	0.557
mean				0.22	0.58	0.2
item1	100% Stacked Bar Chart(4)	Retrieve Value	1.0	0.8	1.0	0.507
item2		Find Extremum	0.0	0.0	0.5	0.819
item3		Make Comparisons	0.3	1.0	0.8	0.583
mean				0.433	0.6	0.766
item1	Pie Chart(5)	Retrieve Value	0.8	1.0	1.0	0.693
item2		Find Extremum	1.0	1.0	0.8	0.979
item3		Make Comparisons	1.0	1.0	0.8	0.944
mean				0.933	1	0.866
item1	Histogram(6)	Retrieve Value	0.7	1.0	0.2	0.914
item2		Find Extremum	0.8	0.2	1.0	0.934
item4		Make Comparisons	0.0	0.9	0.8	0.879
mean				0.5	0.7	0.666
item1	Scatterplot(7)	Retrieve Value	0.0	0.1	0.2	0.633
item2		Find Extremum	0.1	0.6	0.9	0.633
item3		Determine Range	0.6	0.5	1.0	0.527
item4		Find Anomalies	0.0	0.0	0.4	0.341
item5		Find Clusters	0.7	0.8	0.0	0.854
item6		Find Correlations/Trends	1.0	1.0	0.9	0.683
item7		Make Comparisons	1.0	1.0	1.0	0.778
mean			0.485	0.571	0.628	0.635
item1	Area Chart(8)	Retrieve Value	0.5	0.0	0.0	0.683
item2		Find Extremum	0.6	0.6	0.8	0.396
item3		Determine Range	0.2	0.6	0.8	0.381
item4		Find Correlations/Trends	1.0	0.9	1.0	0.944
mean				0.575	0.525	0.65
item1	Stacked Area Chart(9)	Retrieve Value	0.7	0.1	0.9	0.201
item2		Retrieve Value	0.7	0.0	0.2	0.296
item3		Find Extremum	1.0	0.0	0.1	0.909
item4		Find Correlations/Trends	0.7	0.0	0.0	0.417
item5		Make Comparisons	1.0	1.0	1.0	0.959
item6		Make Comparisons	0.8	0.5	0.5	0.286
mean			0.816	0.266	0.45	0.511
item1	Bubble Chart(10)	Retrieve Value	0.0	0.4	1.0	0.864
item2		Find Extremum	0.7	0.1	0.0	0.648
item3		Characterize Distribution	0.0	0.0	0.0	0.457
item4		Find Anomalies	0.7	0.4	0.0	0.628
item5		Find Clusters	1.0	0.9	0.1	0.648
item6		Find Correlations/Trends	0.1	0.2	0.3	0.331
item7		Make Comparisons	1.0	0.2	0.0	0.532
mean			0.5	0.314	0.2	0.586
item1	Choropleth Map(11)	Retrieve Value	0.3	0.1	0.0	0.467
item2		Find Extremum	1.0	0.0	1.0	0.507
item3		Make Comparisons	1.0	0.8	0.0	0.889
mean				0.766	0.3	0.33
item1	Treemap(12)	Retrieve Value	1.0	1.0	1.0	0.934
item2		Make Comparisons	1.0	0.9	0.5	0.879
item3		Identify the Hierarchical Structure	1.0	1.0	0.9	0.879
mean				1	0.967	0.799

Table 3: The table displays the performance of LLM and humans over the VLAT datasets. The table highlights the best performance in each question.

REFERENCES

- [1] B. Alper, N. H. Riche, F. Chevalier, J. Boy, and M. Sezgin. Visualization literacy at elementary school. In *Proceedings of the 2017 CHI conference on human factors in computing systems*, pp. 5485–5497, 2017. 2
- [2] G. Andrienko, N. Andrienko, and A. Savinov. Choropleth maps: classification revisited. In *Proceedings ica*, pp. 1209–1219, 2001. 3
- [3] Anthropic. Introducing the next generation of claude, 2024. Accessed: 2024-06-19. 2
- [4] K. Börner, A. Maltese, R. N. Balliet, and J. Heimlich. Investigating aspects of data visualization literacy using 20 information visualizations and 273 science museum visitors. *Information Visualization*, 15(3):198–213, 2016. 2
- [5] J. S. Bowers, G. Malhotra, M. Dujmović, M. L. Montero, C. Tsvetkov, V. Biscione, G. Puebla, F. Adolfi, J. E. Hummel, R. F. Heaton, et al. Deep problems with neural network models of human vision. *Behavioral and Brain Sciences*, 46:e385, 2023. 2
- [6] J. Boy, R. A. Rensink, E. Bertini, and J.-D. Fekete. A principled way of assessing visualization literacy. *IEEE transactions on visualization and computer graphics*, 20(12):1963–1972, 2014. 2
- [7] Z. Chen, C. Zhang, Q. Wang, J. Troidl, S. Warchol, J. Beyer, N. Gehlenborg, and H. Pfister. Beyond generating code: Evaluating gpt on a data visualization course. *arXiv preprint arXiv:2306.02914*, 2023. 2, 5
- [8] W. S. Cleveland and R. McGill. Graphical perception: Theory, experimentation, and application to the development of graphical methods. *Journal of the American statistical association*, 79(387):531–554, 1984. 2
- [9] E. Firat, A. Denisova, and R. Laramee. Treemap literacy: A classroom-based investigation. In *Eurographics Proceedings*, 2020. 2
- [10] L. Giovannangeli, R. Bourqui, R. Giot, and D. Auber. Toward automatic comparison of visualization techniques: Application to graph visualization. *Visual Informatics*, 4(2):86–98, 2020. 2, 5
- [11] D. Haehn, J. Tompkin, and H. Pfister. Evaluating ‘graphical perception’ with cnns. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):641–650, 2018. 2
- [12] J. Heer and M. Bostock. Crowdsourcing graphical perception: using mechanical turk to assess visualization design. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pp. 203–212, 2010. 6
- [13] A. Inselberg and B. Dimsdale. Parallel coordinates: a tool for visualizing multi-dimensional geometry. In *Proceedings of the first IEEE conference on visualization: visualization90*, pp. 361–378. IEEE, 1990. 6
- [14] S. E. Kahou, V. Michalski, A. Atkinson, Á. Kádár, A. Trischler, and Y. Bengio. Figureqa: An annotated figure dataset for visual reasoning. *arXiv preprint arXiv:1710.07300*, 2017. 1
- [15] Y. LeCun, Y. Bengio, et al. Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*, 3361(10):1995, 1995. 2
- [16] S. Lee, S.-H. Kim, and B. C. Kwon. Vlat: Development of a visualization literacy assessment test. *IEEE transactions on visualization and computer graphics*, 23(1):551–560, 2016. 1, 2, 3
- [17] G. Li, X. Wang, G. Aodeng, S. Zheng, Y. Zhang, C. Ou, S. Wang, and C. H. Liu. Visualization generation with large language models: An evaluation. *arXiv preprint arXiv:2401.11255*, 2024. 2
- [18] H. Liu, C. Li, Q. Wu, and Y. J. Lee. Visual instruction tuning. In *NeurIPS*, 2023. 2
- [19] S. Liu, H. Miao, Z. Li, M. Olson, V. Pascucci, and P.-T. Bremer. Ava: Towards autonomous visualization agents through visual perception-driven decision-making. *arXiv preprint arXiv:2312.04494*, 2023. 1, 2
- [20] A. Masry, D. X. Long, J. Q. Tan, S. Joty, and E. Hoque. Chartqa: A benchmark for question answering about charts with visual and logical reasoning. *arXiv preprint arXiv:2203.10244*, 2022. 1
- [21] M. Mathew, V. Bagal, R. Tito, D. Karatzas, E. Valveny, and C. Jawahar. Infographicvqa. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 1697–1706, 2022. 1
- [22] C. M. McColeman, F. Yang, T. F. Brady, and S. Franconeri. Rethinking the ranks of visual channels. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):707–717, 2021. 6
- [23] K. Mukherjee, T. Rogers, and K. Schloss. Estimating human color-concept associations from multimodal language models. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, vol. 46, 2024. 4
- [24] C. Nobre, K. Zhu, E. Mörth, H. Pfister, and J. Beyer. Reading between the pixels: Investigating the barriers to visualization literacy. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pp. 1–17, 2024. 1, 2, 3
- [25] S. Pandey and A. Ottley. Mini-vlat: A short and effective measure of visualization literacy. *arXiv preprint arXiv:2304.07905*, 2023. 1, 2, 3
- [26] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021. 2
- [27] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022. 2
- [28] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10684–10695, 2022. 2
- [29] P.-P. Vázquez. Are llms ready for visualization? *arXiv preprint arXiv:2403.06158*, 2024. 2
- [30] W. Wang, Q. Lv, W. Yu, W. Hong, J. Qi, Y. Wang, J. Ji, Z. Yang, L. Zhao, X. Song, et al. CogVLM: Visual expert for pretrained language models. *arXiv preprint arXiv:2311.03079*, 2023. 2
- [31] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022. 3, 6
- [32] Z. Xu and E. Wall. Exploring the capability of llms in performing low-level visual analytic tasks on svg data visualizations. *ArXiv*, abs/2404.19097, 2024. 2
- [33] F. Yang, Y. Ma, L. Harrison, J. Tompkin, and D. H. Laidlaw. 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*, pp. 1–17, 2023. 2
- [34] W. Yang, M. Liu, Z. Wang, and S. Liu. Foundation models meet visualizations: Challenges and opportunities. *Computational Visual Media*, pp. 1–26, 2024. 2
- [35] Z. Yang, L. Li, K. Lin, J. Wang, C.-C. Lin, Z. Liu, and L. Wang. The dawn of llms: Preliminary explorations with gpt-4v (ision). *arXiv preprint arXiv:2309.17421*, 9, 2023. 1, 2
- [36] S. Yin, C. Fu, S. Zhao, K. Li, X. Sun, T. Xu, and E. Chen. A survey on multimodal large language models. *arXiv preprint arXiv:2306.13549*, 2023. 2
- [37] J. Yu, Z. Wang, V. Vasudevan, L. Yeung, M. Seyedhosseini, and Y. Wu. Coca: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917*, 2022. 2
- [38] P. Zhang, C. Li, and C. Wang. Viscode: Embedding information in visualization images using encoder-decoder network. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):326–336, 2020. 2
- [39] J. Zong, C. Lee, A. Lundgard, J. Jang, D. Hajas, and A. Satyanarayan. Rich screen reader experiences for accessible data visualization. In *Computer Graphics Forum*, vol. 41, pp. 15–27. Wiley Online Library, 2022. 2