Do Language Model Agents Align with Humans in Rating Visualizations? An Empirical Study

Zekai Shao, Yi Shan, Yixuan He, Yuxuan Yao, and Junhong Wang, *Fudan University, Shanghai, 200433, China* Xiaolong (Luke) Zhang, *Pennsylvania State University, PA, 16601, USA* Yu Zhang, *University of Oxford, Oxford, OX1 3AQ, UK* Siming Chen, *Fudan University, Shanghai, 200433, China*

Abstract—Large language models (LLMs) show potential in understanding visualizations and may capture design knowledge. However, their ability to predict human feedback remains unclear. To explore this, we conduct three studies evaluating the alignment between agent-based agents and human ratings in visualization tasks. The first study replicates a human-subject study, showing promising agent performance in human-like reasoning and rating, and informing further experiments. The second study simulates six prior studies using agents and finds that alignment correlates with experts' pre-experiment confidence. The third study tests enhancement techniques like input preprocessing and knowledge injection, revealing limitations in robustness and potential bias. These findings suggest that LLM-based agents can simulate human ratings when guided by high-confidence hypotheses from expert evaluators. We also demonstrate the usage scenario in rapid prototyping study designs and discuss future directions. We note that simulation may only serve as complements and cannot replace user studies.

arge language models (LLMs) have demonstrated impressive capabilities in encoding visual knowledge and understanding visualizations. They are integrated into the pipelines of visualization generation and recommendation, and function as agents for editing visualizations [1]. A fundamental question is how effectively language models understand visualization, in terms of alignment with human interpretations. Through answering this question, we may gain new insights into visualization perception research [2] and may develop techniques of using language models as low-cost proxies for human subjects in certain scenarios.

Recent works have explored the alignment of lan-

¹Zekai Shao and Yi Shan contribute equally to this work. ²Siming Chen and Yu Zhang are corresponding authors. guage model-based agents with human interpretation in visualization tasks. Some investigate how well model response aligns with human responses in literacy tests by chart question answering [3]. However, the alignments between agents and humans in empirical studies for usability tests or perception tasks are underexplored. A recent study [5] compared LLM-generated and human-made chart takeaways on bar charts with varying layouts to examine LLMs' sensitivity to visualization design choices. However, their study focused on sensitivity to layout as a perception task and takeaway as a response type, leaving other aspects unexplored.

In this paper, we empirically investigate LLM agents' ability to provide feedback on extensive visualization tasks. Ratings, as a common feedback mechanism in empirical studies, are often used to reflect user perception and are prevalent in limited public empirical datasets. Comparing the alignment between agent ratings and human responses in extensive visualization tasks may help identify agent-based

XXXX-XXX © 2021 IEEE

Digital Object Identifier 10.1109/XXX.0000.0000000

evaluation schemes to complement human rating experiments. Therefore, we design a three-stage study on simulating human ratings in visualization tasks to answer the following research questions (**RQs**).

- RQ1: What are the characteristics of agent ratings?
- RQ2: How well do agent and human ratings align?
- **RQ3**: How to better align agent and human ratings?

In the first study for **RQ1** that reproduced a humansubject experiment on time series visualization [6] by GPT-4V [7] we found that the model could simulate human reasoning to give human-like rating values but failed to simulate diverse user profiles. These findings guided the following study design.

In the second study for **RQ2**, we carefully selected five public studies that provided open-source data on the Open Science Framework (OSF)¹ and replicated the original experimental procedures to understand the alignment of agent ratings with human responses. We analyzed the alignment of agent ratings with human responses in the replication of six experiments. By also inviting five experts in visualization evaluation to code confidence for individual experiments, we suggest that the alignment of agent ratings with human responses may positively correlate with expert pre-experiment confidence. Agents, as aggregations of historically trained knowledge, can mimic human responses but cannot replace human subjects.

In the third study for **RQ3**, we conducted four additional experiments and observed interesting outcomes. We found that automatic external knowledge injection may enhance the alignment of agent ratings with human ones. However, emphasizing variables intended for comparison in prompts or presenting textonly inputs without the corresponding visualizations might lead to obvious biases from original ratings.

Additionally, we proposed an experimental scenario that involved both human-subject and agent-based studies, illustrating agents' potential in rapidly prototyping experimental designs.

Our contributions are summarized as follows:

- We design three studies based on existing experiments with human subjects to illustrate the relative alignment between agent ratings and human responses, demonstrate how this alignment correlates with experts' confidence before experiments, and explore the potential for further enhancing alignment through constructive strategies.
- We demonstrate the potential scenarios for agents in swiftly evaluating prototypes and discuss future directions for simulating visualization experiments.

The examples of prompts and experimental codes can be found in the appendix and the complete codes are available in https://github.com/ZekaiShao25/ Agents-Ratings-in-VIS-Experiments.

Related Work

We review the related work in visualization evaluation, large model and agents, and machine understanding of visualizations.

Visualization Evaluation

User evaluation is a long-standing topic in visualization research. Many works categorize evaluations in terms of purpose, subject, and scenario. From the perspective of collected data types, empirical study data includes quantitative data, such as accuracy, task completion time, and questionnaire ratings, as well as qualitative data, such as think-aloud results and interview records. We focus our investigation on the alignment of ratings simulated by large model-based agents with human data.

Empirical studies in visualization research span many topics, such as dimension reduction algorithms, network representations, time series [6], [8], parallel coordinates, aesthetics [9], uncertainty [11], and visualization literacy. Additionally, some work has assessed interactive visualization systems with case studies and interaction log analysis. We limit the scope of our study to rate basic 2D visualizations without interaction.

Large Model and Agents

LLMs, such as GPT-4 (no-vision) [7] and Llama, have shown exceptional abilities. Recently, LMMs, such as GPT-4V and GPT-4o, have exhibited visual understanding abilities and general intelligence, employed as autonomous agents for general tasks such as web navigation and interpreting webpage screenshots. To address hallucination issues, retrieval-augmented generation (RAG) ensures accuracy and relevance, particularly in fields like scientific research. Additionally, crowdsourcing-inspired workflows have been proposed to optimize LLM chains for better performance. These studies motivated us to investigate large models in simulating human ratings and inspired the design of our agent-based procedure.

LLM-based modeling is also being applied to visualization generation and recommendation, natural language dataset synthesis, enhancing visual analytics systems, and visual storytelling. Besides integrating LLMs into automatic pipelines, Liu et al. [1] utilized

¹https://osf.io/

LMMs as autonomous agents for visualization modification. In this work, we study the alignment between large model-based agents and human ratings.

Machine Understanding of Visualizations

As machine learning advances, visualization community has adopted models to enhance the understanding of visualizations. Previous research has employed information theory and visual saliency automatic quality evaluation or explored models like CNNs for visualization perception tasks [2], [16], [17]. For example, Haehn et al. [17] replicated Cleveland and McGill's classic perceptual tasks across five visualization encodings using multiple CNNs, demonstrating that while CNNs can match or exceed human performance on simple cases, they often fail on more complex visual encodings, revealing their limitations as models of human graphical perception. Our study shares a similar goal of modeling human visual perception but explores whether more capable multimodal LLMs can simulate human judgments in rating tasks.

With LLMs' emergence, researchers have extended these studies to evaluate their capabilities in visualization tasks that align with human-subjective study settings, including low-level analytics [13], literacy assessments [3], misleading visualization detection [14], and chart takeaway generation [5]. These studies often investigate parameters and settings such as different prompts and models. They can be categorized into two primary approaches: one focusing on task automation, aiming to maximize LLM accuracy [13], [14], and the other also measuring alignment between LLM outputs and human responses [3], [5]. Beyond direct promptbased models, Wang et al. [15] proposed DracoGPT by using Draco to extract design preferences of LLMs, and Cui et al. [18] generated datasets for visualization literacy tests with LLMs.

Our work aligns with the contemporaneous research on human-LLM alignment on visualization perception tasks but centers on the relatively underexplored area of subjective ratings as a feedback type. Additionally, our empirical study spans multiple public empirical studies involving subjective ratings. We also explore the potential of RAG-enhanced agents to improve alignments and address future application scenarios.

Study I: What Are the Characteristics of Agent Ratings?

We first conducted Study I to replicate an experiment reported in Adnan et al.'s work published in ACM CHI 2016 [6]. Our Study I explored the alignment between agent-simulated ratings and the ratings from human subjects. We prioritized three factors for selecting the existing experiment: ease of experimental replication, categorization as 2D visualization perception tasks without interaction, and availability of collected and analyzed human ratings. We chose to replicate the results on the improvement of user experience by a visualization design for time series data.

Original Study

Study design. The original study employed a withinsubjects design for all independent variables: 4 interaction scenarios (no interaction, highlighting, tooltips, and highlighting-tooltips), 3 visual encodings (position, color, and area), 2 coordinate systems (Cartesian and Polar) and 4 tasks (maxima, minima, comparison, and trend detection). The time series visualizations presented to the participants are created from synthetic time series datasets. The maxima and minima tasks are to identify the highest and lowest value in the visualization, respectively; the comparison task is to compare the aggregated value of two slices of the time series; and the trend detection task is to identify a subset of the data with the smallest upward trend.

Original results. The authors analyzed the outcomes of completion time, accuracy, confidence level, and ease of use and obtained some conclusions (**C**). We focus on their conclusions on the impact of visual encodings on ratings, which are listed as follows. The area encoding received lower scores in other tasks (**C1.1**) but excelled in comparison tasks (**C1.2**).

Agent-Based Study

Our agent-based study excluded one independent variable, interaction, and one value of another variable, the polar coordinates, in the original design, due to the impracticality of enabling the agent to perform interactions and understand uncommon visualizations. We focused on having the agent work with line charts, heatmaps, and icicles, and evaluating agent-simulated ratings of ease-of-use and confidence level. We further explored using GPT-4 to simulate rating for pixel-based visualization tasks, emphasizing the alignment with the results on ease-of-use and confidence levels by human subjects in the original study.

Our goal is to align the agent's ratings with established results (**C1.1, C1.2**) and also offer human-like reasoning in its explanations. To achieve this goal and utilize its internal knowledge, we structured the task completion in three steps. We prompted the agent to output its internal knowledge about the given visualization designs, which is incorporated into the input of the next step. Then, we presented the task description and visualization images, prompting the agent to outline the steps for task completion and potential challenges. Next, we bundled the output from two steps and rating criteria as the final step's input, and asked the agent to assign ratings. We experimented with one and three visualizations per input, respectively. We observed that ratings can be identical when given one image per session, but show clear differences when collectively given three at a time, with the latter results being more reasonable. This observation shows that agents differ from humans as they apply inconsistent standards for different experiment settings across sessions. This finding also affected our subsequent procedure design. Thus, we input three visualizations together to GPT-4V for each task with 30 repeated trials.



FIGURE 1. Agent ratings for time series visualization tasks.

The agent's ratings, as shown in Fig. 1, indicate that for tasks other than comparison, agents assigned significantly lower scores to area encoding, aligning with **C1.1**. For instance, GPT-4V describes that identifying maxima tasks is most efficiently done by locating the highest point on a line chart. Heatmaps rank second due to their graduated color mapping, while the icicle plot presents the greatest challenge due to the narrow width of each block and non-adjacent comparison. However, for comparison tasks, agents also concluded similarly, which contradicts **C1.2**. It overemphasized aggregating daily-level blocks for weekly sales calculations, overlooking the possibility of directly comparing weekly sales between two middle tiers.

We noted GPT-4V's knowledge in step 1 that "for icicle charts, child nodes are added below their respective parents." This implies that the unchanged outcomes may stem from the added knowledge not clearly explaining the quantitative relationship between parent and child nodes. Thus, we replace the selfprompted internal knowledge with "For the icicle plot, to aggregate child nodes, consider the size of the parent node", as the partial input of step 2. As anticipated, this adjustment led GPT-4V to propose logical completion steps and rate the icicle plot's confidence to the highest level. Given these insights, we **remain open to the**

idea of employing agents for subjective rating.

We then aimed to assess whether the agent could embody various roles to mimic diversity within or across groups. Based on the original study's 24 participant descriptions, we created 24 distinct user profiles varying in gender, age, identity, and education, incorporating these user profiles into the first step's input. Yet, varying profiles scarcely affected the outcomes, suggesting GPT-4V's limited grasp on how different demographics influence user behavior, thus failing to represent individual differences. Additionally, we adjusted the agent's internal parameter (temperature) to increase the diversity of responses, yet the inherent randomness of the model could not accurately simulate user diversity. We conclude that while agents may simulate the behaviors of an "average user" by capturing common patterns learned from large-scale training data, it is still challenging to produce data mimicking diverse users in visualization tasks.

Study II: How Well Do Agent and Human Ratings Align?

We conducted experiments to investigate the alignment between agent ratings and human responses align by comparing whether they lead to consistent conclusions in specific tasks (**RQ2**).

Data Preparation

We gathered a wide array of visualization evaluation literature. Similar to constraints in Study I, we focused on public studies with complete original user study materials. We limited the scope of papers to prominent conferences and journals: VIS, TVCG, EuroVis, CGF, and CHI. The criteria and process for the inclusion of literature can be concluded as follows:

- Keyword matching. We first applied keywordbased regular matching, requiring titles or abstracts to include terms like "evaluate," "investigate," "user," "viewer," "crowdsource," "assess," or their variations.
- OSF filtering. We then filtered the matched papers for those with experimental materials on the OSF to reduce replication costs and increase the reliability of our evaluation dataset.
- 3) Manual filtering. At last, we manually reviewed all filtered papers and excluded papers not related to visualization evaluation, which are surveys, those evaluating newly proposed visualizations, and those involving interaction or 3D representations.

The results of each step are shown in (Fig. 2). In summary, we identified five suitable papers meeting our requirements: one in CHI (Timeline [8]), one in



FIGURE 2. The data preparation steps. The values represent the number of remaining papers after each step. For VIS papers, we skip the keyword matching step and utilize Vispubdata dataset to obtain all paper titles.

VIS (Imputation for Uncertainty [11]), three in TVCG (Fit Estimation [10], Texture [9] and Magnitude Judgement [12]), and none in EuroVis or CGF. The five papers form the experimental datasets.

Our study primarily utilized GPT-4V (temperature = 0) as the rating agent, the most powerful multimodal LLM at the time of conducting this project. We will discuss the significance of our work for state-of-the-art LLMs in the discussion. We also tested other LMMs across varying model sizes and types, with a goal to deepen our understanding of how model size and type influence the results.

General Procedure

We identified the variables involved in each experiment, such as visualizations, tasks, and data. We focused on simulating conclusions related to rating, which include participants' confidence level, aesthetics, and readability topics.

Replicating Original Study with Prompting Agents To ensure the integrity of our experiments, we strive to replicate the original experiment procedure based on the open-source data as closely as possible within the constraints presented by the current capabilities of the agent. This replication involves inputting training tutorials and task descriptions into the prompt without iterative prompt optimization, using the original study materials directly to provide a contextual foundation for agents' operations. If rating is the task itself, the agent assigns points in a single operation. If the original study required completing the task first before ratings, then the agent proceeds in a two-operation chain: first outlining steps and difficulties to complete the task, then assigning ratings based on those steps. However, we had to modify the experimental setup in two aspects.

Transitions from between-subjects to withinsubjects variables. The necessity for such modifications arises from the lack of standardized rating criteria for comparing visualizations. They result in the agent's inability to maintain a consistent standard across time and images as humans do, as shown in Study I: What Are the Characteristics of Agent Ratings?. As a result of our modification towards within-subjects variables, different sessions shared an identical experiment setting, making the results meaningful.

Transitions from global comparison to local comparison. The agent's token limit restricts image input, with GPT-4V's API often failing beyond nine normal-resolution images. Furthermore, it's not feasible to fix a specific image as a standard and use its rating as a benchmark across multiple sessions for comparison (see observations in Study I: What Are the Characteristics of Agent Ratings?). Hence we had to limit image batches by selecting and inputting the most relevant images.

Potential Limitations for procedures of our agent-based study. We shuffled the input images within one session to eliminate the learning or fatigue effect brought by our modification to within-subjects design, a practice frequently adopted in human-subject studies. However, due to the input limit, shuffling a small set of images may not mitigate the effect as effectively as a larger set does in human studies, since images are not as disordered as in a larger shuffled set. Moreover, our sample size was insufficient to form a nearly uniform distribution of permutations to counterbalance the order effect. For experiments where conclusions can be drawn by comparing images within the input limit, we attempt to replicate their findings but not all experimental phenomena. Conversely, current agent-based studies are not feasible for replication if the quantity of inputted images is insufficient for the conclusions. We acknowledge these limitations above and will discuss how potential flexible design can alleviate them.

Results Analysis by comparing alignments of conclusion Generally, we focused on the alignment between the agent's feedback and original conclusions (C) and hypotheses (H) (predefined or not), rather than data distributions. Our observations show that agents' baseline of rating differs from human evaluators. This is especially evident when agents rate on broader scales such as 0-100, where their average ratings may endogenously higher or lower than human participants' ratings, and the differences should not be ignored. For example, in Fit Estimation study [10], the ratings of human raters are mainly distributed between 45-90 in the 0-100 range, while the ratings of agent raters are mainly distributed between 70-85 (Fig. 3). In our analyses, the problem can be resolved by comparing conclusions, instead of referring to absolute values. In addition, our analyses were based only on sample information, where we employ empirical distribution.

Experiment I: Fit Estimation

Original Study **Materials and Procedure.** Reimann et al. [10] aimed to understand the patterns in visually estimating the fit between models and data. Two main experiments were conducted involving 62 German students to rate how well the data points fit the shown line.

The first experiment delved into the relationship between the noise level from the model and its visual perception of fit ratings. It involved participants using an online interface to sequentially rate the fit for 30 scatterplots, presented in random order, based on a 5 (noise levels) x 6 (parallel versions with random values) within the subject. Using a slider bar, participants rated the fit on a scale from 0 (very low fit) to 100 (very high fit), with instructions and visualization provided. Each generated scatterplot featured a consistent model line and varied noise levels were achieved through adjusted y-values for each x-value. We identified with two hypotheses. The first is a significant influence of noise on visual estimations (H2.1). The second is a negatively accelerated relationship between noise levels (statistical deviations) and visual perception (H2.2), derived from Fechner's law of psychophysics.

The second experiment explored the influence of the decentering of data points from the model on fit estimation. It replicated the first one's methodology, presenting participants with randomized 42 scatterplots based on a 6 (noise levels) x 7 (decentering conditions) within-subjects design. Each scatterplot was manipulated with noise by varying the standard deviation (sd) from 0.2 to 1.2 in steps of 0.2. Decentering was adjusted by changing the mean y-value with a certain ratio for each x-value, with seven conditions ranging from neutral (ratio = 0) to three increasing levels (ratio = 0.25, 0.5, 0.75) of decentering both upwards and downwards the model line. Evaluators expected that increased decentering would correlate with lower fit estimations (H2.3) at constant noise levels. However, they were uncertain if participants would strongly weigh decentering, leading to no predefined hypothesis.

Results. For the first experiment, a repeated measures ANOVA confirms that noise significantly affects visual fit estimation ($F(4, 244) = 273.04, p < .001, \eta_p^2 = .82$), with all pairwise comparisons significant (Bonferroni corrected) (**C2.1**). Visual deviation ratings increase more sharply at lower noise levels compared to higher ones, with *t*-test confirming these differences in rate changes as significant (**C2.2**), t(61) = 7.61, p < .001, d = 0.97. These conclusions from the user study align with both two hypotheses.

For the second experiment, an ANOVA analysis combined upward and downward decentering, shows significant effects of both decentering, $F(1.43, 87.23) = 28.25, p < .001, \eta_p^2 = .32$ (**C2.3**), and noise on fit estimations without a significant interaction between these factors. Thus, **C2.3** is fully aligned with **H2.3**. Decentering has a robust but surprisingly small impact compared to noise (**C2.4**), as the largest difference in fit estimations (M = 36.27, SD = 14.19) between the highest and lowest noise levels, is significantly greater than the difference between centered and strongest decentered conditions (M = 8.48, SD = 11.39), t(61) = 11.32, p < .001, d = 2.16.

Agent-based Study Material and Procedure. For the first experiment, we verified C2.1 and C2.2 through 6 sets of parallel versions, where GPT-4V evaluated 5 scatterplots at varied noise levels, repeated 5 times to minimize randomness. For the visual input, all scatterplots were generated using the open-source code that generates stimuli in OSF. The order of the inputted five scatterplots was random, which is consistent with the original study. For the textual prompt, we began with the role-play instruction "You are an average user" and concluded with "Please list the ratings and reasons in the order of the images." The remaining prompts were directly sourced from the tutorial and task descriptions in the original paper and OSF.

The original second experiment further explored C2.3 and C2.4 totaling 42 images, which exceed the input limits as stated in Replicating Original Study with Prompting Agents. Therefore, we performed a local comparison with three methods: comparison of decentering across four levels against a fixed noise level, comparison of decentering effects at noise levels of 0.2 and 1.2, and comparison in specific conditions (e.g., strong decentering at noise=0.2 vs. neutral decentering at noise=0.4). The first comparison aimed to simulate C2.3, the second one targeted C2.3 and C2.4 verification, and the last emerged from the close ratings of pairs in the original experiment. In the first comparison, the agent rated 7 scatterplots with varying decentering but fixed noise levels, repeating 5 times. In the second comparison, the agent rated 8 scatterplots each time, including 2 with no decentering and another random (upwards or downwards) 6 with different decentering for 50 parallel trials. In the third comparison, the agent rated 3 scatterplots from each condition, repeating 5 times. The prompts were the same as in the first experiment.

Results. As we decided to make no assumptions on the agent's output distribution across sessions, we chose not to use ANOVA or *t*-tests. For each condition,



FIGURE 3. Results of agent ratings for fit estimations under various noise levels (top) and decentering (bottom). Error bars represent 95% bootstrap Cls.

we employed bootstrap (10,000 replicates) to estimate the 95% CI of the means. The results of the first agentbased experiment are shown at the top of Fig. 3. Agent ratings largely confirm **C2.1**, but do not support **C2.2**. This is because the scores show minimal change from noise levels 0.5 to 1, and experience a significant drop from 1 to 1.5, with the non-overlapping CIs indicating a significant difference between these means.

Agent-based ratings for the comparison in the second experiment are presented at the bottom of Fig. 3. It suggests that the agent ratings support C2.4, as the influence of noise substantially outweighs that of decentering. The results partially support C2.3, showing that at a constant noise, strong decentering results in lower ratings compared to no decentering. However, fit ratings do not consistently drop with increasing decentering, suggesting the agent perceives decentering less acutely than humans. The results of other comparisons partially align with the user study, which can be found in our codes repository. Additionally, the score with sd=1.2, displayed at the bottom of Fig. 3, is significantly lower than the score for sd=1 shown above, again showing that the agent's analytical perspective is not consistent in different experiment settings. This result indicates the necessity of our modification from between-subjects to within-subjects variables.

Experiment II: Imputation for Uncertainty

Original Study **Materials and Procedures.** Sarma et al. [11] explored how different representations and uncertainty affect chart interpretation. They used a mixed design with six representation types (no imputation, mean, CI, density, gradient, and HOPs with animation) as between-subjects treatments, and two tasks (estimating average and trend) and two proportion levels as within-subjects treatments. The tutorial explained the

representation, missing values, and their proportion. It stated that missing values were estimated using an appropriate imputation method, noting the inherent uncertainty of these estimates. Participants also performed a practice task with feedback. Regards ratingsrelated results, participants' confidence was gathered using a 5-item Likert-style question on the interface.

Results. They applied a Bayesian hierarchical ordinal regression model to estimate participants' response probabilities on Likert items, analyzing the effects of representation, proportion, trial_id, block_id, and participant id. Trial and block items were used to model learning effects, while they were marginalized as no meaningful effects were found. The last item was used to model individual characteristics. For ratings, they computed the likelihood of an average participant rating their confidence as 3 or higher on the Likert scale, Pr(>3), across representation conditions. Results showed that participants had the highest confidence with no imputation in average estimation Pr(> 3) = 64%, slightly less with mean imputation Pr(>3) = 52%, and even lower in conditions showing uncertainty (C3.1). In addition, there were no significant differences in confidence in trend estimation (C3.2). They made the hypothesis that presenting mean imputation estimates alone would result in the highest participant confidence, regardless of task types (H3.1, H3.2). They expected that no imputations would reduce confidence due to a lack of information, with participants aware of this fact. They expected that explicitly showing uncertainty would diminish confidence due to the direct acknowledgment of uncertainty. Thus, H3.1 partially aligns with C3.1, while H3.2 is disproven.

Agent-based Study Material and Procedure. Based on the above description, we shifted to a withinsubjects design for variables of representation. We excluded HOP as it contains animation, and the agent couldn't process it. For one task (estimating mean or trend), the agent completed 2 (proportion) x 2 (error) x 2 (condition) = 8 trials for the inputted five representations. Each trial was repeated 5 times to reduce randomness, resulting in 40 sessions. For visual input, we used original images from the OSF. The textual prompt comprised direct copies of descriptions from the tutorial and interface, with the added role-play instruction, "You are an average user." Following the procedure where participants first complete the task and then provide ratings, the agent first generated steps based on the task description and visual images, subsequently, rating based on these steps.

Results. In our initial model selection process, we supposed the original regression model was imprac-

tical due to the lack of prior knowledge about the distribution or assumptions of the agent's output and other potential factors. We discarded variables like participant and learning factors due to our modifications to the experiment procedure and tried adding random intercepts for each session. However, the regression model did not capture data well, as posteriors were largely different from the empirical distribution. Therefore, we employed descriptive statistics to analyze the frequency distribution of agent ratings for each imputation method, by comparing the proportion of agent ratings \geq 3. Fig. 4A reveals that in average estimation, the agent rated mean imputation highest $(P(\geq 3) = 57\%)$, with baseline $(P(\geq 3) = 48\%)$ and gradient ($P(\geq 3) = 50\%$) methods scoring lower, aligning with H2.1 but not with C2.1. In trend estimation, mean, baseline, and CI received high scores with the same proportion, against both C3.2 and H3.2 as other methods scored significantly lower.

Summary of Other Replication Experiments

We carried out four additional replication experiments on GPT-4V. Among 9 conclusions in these experiments, the agent's ratings were only aligned with **C4.1** in Timeline (Tab. 1). We also conducted tests on opensourced models including LLaVA-series model (7B and 13B) and BakLLaVA 1 (7B). Due to the space limit, we briefly summarize our findings here. Detailed study design and results are available in the appendix.

- Agent ratings may align with general human preferences, though it falls short of simulating detailed findings. In Timeline [8] experiment, the agent consistently preferred horizontal lines for any dataset or task type as the most readable timeline visualization.
- Agent may not simulate aesthetic and textural readability. In Texture [9] experiment, agent ratings cannot align with any original conclusion related to texture.
- Agent may focus on textual elements within images. In Magnitude Judgement [12] experiment, the agent always offered an anti-human interpretation of axis manipulation by focusing on the text labels.
- Open-sourced multimodal agents demonstrate poor comprehension in visualization tasks.

Confidence Coding and Implications

Confidence Coding. Agent ratings usually align with human data for basic visualization tasks that seem predictable (C1.1, C2.1) and often resemble expert hypotheses (H3.1). This inspires us to consider the relationship between their alignment and the confi-

dence level of human experts in forming hypotheses. Since the original study's expert evaluators typically did not emphasize hypothesis confidence, we invited and consulted five external visualization experts for confidence coding. All experts have had over three years of experience in visualization research, designed and conducted formal user studies, and have more than three publications in TVCG and VIS. They were not our co-authors and had not read the papers being evaluated. We engaged in one-on-one interviews with each expert for 40 minutes, using slide sharing to present the interview materials. We first introduced the interview process and the hypothesis confidence scale using one slide. Each subsequent slide presented one user task and related experimental materials and any relevant background information mentioned in the study (i.e., relevant research findings or theories from other fields), but excluding the original evaluators' hypotheses. Experts were asked to verbally decide on forming their hypotheses for each experiment and give confidence in three-scale Likert ratings. If they could form a hypothesis, they would assign a high (3) or medium (2) confidence level, indicating strong or moderate certainty in their predictions about general user performance, respectively. If they were unable to form a hypothesis, they assigned low (1) confidence.

For cases with hypotheses in the original study, we found that the hypotheses from five experts largely matched. For the few opposing hypotheses (no more than one in five), we lowered their confidence level to low. In cases without original hypotheses, external experts' hypotheses were more diverse. We adjusted the confidence levels to low for experts who proposed opposing hypotheses. This change highlights that experts do not have unified assumptions. We averaged the modified confidence levels of the five experts for each conclusion, rounding up to determine human expert confidence, as illustrated in the "Confidence" column of Tab. 1.

Implications. The results from our agent-based replication experiments (black rows in Tab. 1) can be summarized as follows:

- Ratings from the agent often align with human results for basic tasks where experts have high confidence. In experiments where external experts had high confidence, the match score of agent ratings is 3/5 (C1.1, C2.1, and C4.1); in other cases, the match score is 1/12. These three Cs represent the most basic and general findings of each experiment, involving only a limited number of variables.
- The alignment of agent ratings drops when expert confidence is moderate and relies on ad-

TABLE 1. A summary of our experiments on ratings investigating the alignment among hypothesis (H), conclusion (C) and agents' feedback (A), with "P" denoting a partial alignment. Black rows represent replicate experiments for **RQ1** and **RQ2**, light rows represent additional experiments for **RQ3**.

Experiment	Conclusion	Confidence	Hypothesis	Alignment		
				H-C	H-A	C-A
time series [6]	C1.1	H (2.93)	N	/	/	Y
	C1.1	H (2.93)	N	/	/	Υ
	C1.2	H (2.80)	N	/	/	N
	C1.2	<u>H (2.80)</u>	N	/	/	N
fit estimation [10]	C2.1	H (2.60)	Y	Y	Y	Y
	C2.2	M (2.20)	Y	Y	N	Ν
	C2.3	M (2.20)	Y	Y	Р	Р
	C2.4	L (1.40)	N	/	/	Y
imputation for uncertainty [11]	C3.1	M (1.80)	Y	Р	Y	Р
	C3.1	M (1.80)	Y	P	P	Y
	C3.2	L (1.40)	N	N	N	N
timeline [8]	C4.1	H (3.00)	Y	Y	Υ	Y
	C4.2	M (2.47)	Y	Ν	N	Р
	C4.2	M (2.47)	Y	Ν	Y	N
	C4.3	L (1.42)	N	/	/	Ν
texture [9]	C5.1	L (1.40)	N	/	/	N
	C5.2	L (1.20)	N	/	/	Ν
	C5.3	M (1.80)	Y	Р	N	Р
	C5.4	M (1.90)	Y	Р	Ν	Ν
magnitude	C6.1	H (2.80)	Y	Y	Ν	N
judgement [12]	C6.2	M (1.60)	Ν	/	/	Ν

ditional knowledge or unreliable common sense. In such cases, agent feedback never fully aligns with the conclusions, often showing partial alignment or complete opposition, regardless of original hypothesis-conclusion alignment.

Agent ratings are almost useless when experts have low confidence and set no consistent predefined hypothesis. The rest are low-confidence scenarios usually without original hypotheses, where agent ratings seldom support the original conclusions, except for alignment with C3.4.

Our response to **RQ2** suggests that the alignment of agent ratings with human data might correlate with expert confidence for hypotheses. Agents extract basic information from images, integrate it with task-specific text prompts, and leverage their internal historical knowledge derived from training to approximate human behavior and predict user study outcomes.

Study III: How to Better Align Agent and Human Ratings?

Leveraging the insights gained from simulating the original procedures, we also dedicated our engineering efforts to achieving more aligned results by proposing specific strategies (**RQ3**).

How Do Visual and Textual Cues Influence Reasoning?



FIGURE 4. Proportion of agent ratings \geq 3 for imputation methods addressing uncertainty. They feature two tasks (estimating the average and trend) across three conditions: (A) text and images combined, (B) text only with the same prompt, and (C) text only with a general description.

Prior work shows that multimodal agents exhibit varied sensitivity to textual and visual prompts—text-centric prompts can sometimes enhance reasoning [4]. This motivated us to explore agent-based studies with minimal reliance on images. We began with Time Series to assess the impact of image modifications and further tested Imputation for Uncertainty to examine the role of keywords in text prompts for reasoning.

Time Series: agent ratings rely on general perception and knowledge, but ignore visual details such as color. Color plays a crucial role in human perception for visualization. Although the agent provided human-like ratings in Study I, its explanations lacked specific references to the color encoding of heatmap. Consequently, we altered the visual input to grayscale images while keeping the textual input unchanged. The change yielded similar scoring distributions, with explanations omitting "grayscale" or color-related words.

Imputation for Uncertainty: using only text prompts for general tasks without specific scenarios increases the risk of data pollution. We discarded the visual inputs and modified the textual prompts in two ways: (1) by retaining the original textual prompt (Fig. 4B); and (2) by removing specific details about the data, graphs, and task-for example, simplifying "estimate the average value of X of the whole dataset (including both Cyan and Orange groups)" to "estimate the average value of the whole dataset" (Fig. 4C). It shows that with images removed, the agent assigned more extreme but consistent scores, either above or below 3. Yet, keeping the original textual cues in the prompt resulted in entirely inaccurate scores, like all baseline scores falling below 3 ("Mean" tasks of Fig. 4B). This occurred because textual cues carry significant detail-related information (e.g., groups in scatterplots, color), and the absence of images leads to incomplete inferences by the agent. Conversely, directly inquiring on the study's final concerns yielded the results that are consistent with C3.1, with the average score for the baseline

surpassing the mean (3.85 vs 3.6). The agent noted in its explanation that "baseline will make users feel more reliable, although this may be overconfident," which aligns with the original study's post-experiment analysis. This result raised a concern with data pollution in the discussion.

We also conducted experiments on Timeline. The procedures and results are included in the appendix. We found that explicit comparisons of variables may make the agent rely on the text to offer stereotypes similar to those of experts.

Does Knowledge Injection Improve Alignment?

Experimental Design. The success of manually adding specific knowledge to icicle plots in Study I: What Are the Characteristics of Agent Ratings? and the popularity of RAG technology underscores the potential of automatic knowledge injection in our context. RAG technology involves three steps: 1) compiling a corpus and training a vector database; 2) querying the needed knowledge with natural language; and 3) incorporating the related knowledge into the prompt for answering. Given the absence of a specific corpus for visualization evaluation, we employed web retrieval as our knowledge source. The workflow of automatic knowledge-assisted answering is illustrated in Fig. 5. The initial two steps (Fig. 5A) were the same as those in Study I: What Are the Characteristics of Agent Ratings?, where the agent first described visual graphs and then scheduled steps for task completion. Different from the initial replication experiment, we input specific visualization names, like "icicle plot," into another websearch-enhanced agent (GPT-4) and prompted it to gather related knowledge from web pages, stored in a database. Then we split the agent-simulated steps into sentences, each serving as a query into the websearch-sourced database to retrieve knowledge blocks by semantic similarity of vector embeddings (Fig. 5B). For each sentence query, we repeated 3 trials to identify the most related knowledge block. This top-1 block was then embedded in the schedule prompt to guide subsequent step generation and ensure knowledgeinformed ratings (Fig. 5C).

Experiment Results. Incorporating three visual graphs as the inputs, our RAG-enhanced procedure yielded promising results. While the rating distribution for other tasks remained largely the same, icicle plots in the comparison task saw significant improvement, making their overall ratings closer to the other two visual encodings (compared to Fig. 1). A detailed analysis showed that although icicle plots never ranked



FIGURE 5. RAG process designed for knowledge injection in the experiment of Time Series in Study III.

first in the original procedure in Study I: What Are the Characteristics of Agent Ratings?, they occasionally achieved the highest scores in RAG-enhanced procedure (19.0% for confidence level; 20.7% for ease of use). This improvement can be attributed to the action to correct a previously incorrect step involving icicle plot comparisons. This approach suggests "*aggregating all sizes of the weeks*," while the relevant search found "*where the size of parent nodes directly corresponds to the sum of their children's sizes*." This knowledge guidance effectively alters the steps and ratings consistently.

Implications. This result illustrates the potential help of using external knowledge injection. Yet, we recognize potential issues with this approach. The accuracy and relevance of web search information to the agent's erroneous steps are uncertain, limiting the generalizability of our method. Furthermore, even with improvements from automatic external knowledge guidance, the results seem to need more scrutiny and verification before being used.

Potential Scenarios: Fast Prototyping Study Design



FIGURE 6. Scatterplots: strong decentering (0.99) with 0.2 noise (R1) and 0.4 noise (R3); no decentering with 0.4 noise (R2) and 0.6 noise (R4).

The Goal of Experiment. In the original study of fit estimation [10], strong decentering was described as the data shifting by 0.75 times the noise size, where points were still on both sides of the line as shown in the OSF materials. The results indicate that 1) the ratings for *(strong decentering, 0.2 noise)* (R1) are higher than those for *(no decentering, 0.4 noise)*

(R2); and 2) the scores for (strong decentering, 0.4 noise) (R3) are on par with those for (no decentering, 0.6 noise) (R4). We doubted the effect of decentering and plan to test a more extreme decentering (shift value=0.99) scenario (Fig. 6). It would make points almost clustering on one side of the line as shown in the R1 and R3 of Fig. 6. As outlined in Experiment I: Fit Estimation, agent ratings supported **C2.4** and revealed similarities to human judgments in specific local comparisons. Thus, we aimed to compare the scores from agents for extreme decentering against real user study outcomes to investigate its alignment in swiftly adjusting experimental parameters.

Experiment Design. We conducted a new withinsubject study with a 3 (noise levels: 0.2, 0.4, 0.6) \times 3 (decentering conditions) design, involving 38 participants. One sample of suspicious rating pattern was removed for short completion time and casual ratings across all images. The procedure and instructions replicated those in the original study, using an online interface integrated with guestionnaires and visual stimuli. After providing informed consent, each participant went through 9 scatterplots in a randomized order and rated them via slider bars. This procedure ensures a balanced dataset for ANOVA and covers R1-R4 (Fig. 6). For the agent-based study, due to the input limit, we excluded the three images with (no decentering, 0.2 noise) and (strong decentering, 0.6 noise) from the agent's input, as they were irrelevant to R1-R4. We used GPT-4V to complete the withinsubject ratings for six images with 20 repetitions to mitigate randomness.

Results. We combined the ratings for upward and downward decentering (there was no effect of direction) into R1 and R3 by averaging, as the original study did. For the study involving human subjects, a repeated measures ANOVA analysis revealed a significant main effect of decentering, $F(1,36) = 14.00, p < .001, \eta_p^2 = .32$, and noise, $F(1.6, 57.60) = 97.047, p < .001, \eta_p^2 = .73$ (both Greenhouse-Geisser corrected). No significant interaction effect was found, $F(1.90, 68.22) = .003, p = 1.00, \eta_p^2 = .00$. Graphical inspection indicated only minor deviations from the normality assumption required for ANOVA. These conclusions are consistent with those presented in the original study.

Fig. 7 shows the ratings of agent (left) and humans (right). The results of the new user study align with the original study. Even if decentering reaches a more extreme value of 0.99, the result of R1>R2>R3=R4 is still valid. Agent feedback is also consistent with the two user studies, and the impact of decentering on it is still weaker than that on humans.

Implications. Our experiments showed two poten-



FIGURE 7. Fit ratings by agent (left) and human (right) for different conditions of noise and decentering. Error bars represent 95% CIs.

tial use cases. If agent outputs align with the existing user study, and new studies only adjust experimental parameters without altering other procedures, reflecting these parameter changes in the agent's input may allow for the preliminary simulation of the new studies. In addition, if the results of a small-scale pilot study align with the predictions by the agent, evaluators may adjust parameters and watch for significant output changes to set parameters for the formal study. This approach assists in setting parameters that align with experimental objectives. However, more empirical evidence from studies with large datasets is needed to further confirm this finding.

Discussion

Possible misuses. Novice evaluators may be influenced by the agent's human-like outputs as demonstrated in our study, despite their inherent unreliability. Moreover, verifying qualitative data proves more difficult than text data, possibly amplifying reliability concerns regarding crowdsourced data [19]. Currently, closed-source large models like GPT-4V can rate visualizations in a human-like manner, with high usage costs serving as a temporary barrier. However, rapid advancements in open-source models may soon reduce these barriers, increasing associated risks.

Keeping Pace with State-of-the-Art LLMs. More advanced LLMs are emerging that could produce ratings more aligned with human responses. We reproduced the agent-based study in Fit Estimation using GPT-40 and presented the results in the supplementary material. While GPT-40 differs from GPT-4V in rating range, its alignment with humans mirrors that of GPT-4V. We believe our study methods and opensource code can be directly applied to any LLM, and encourage others to use them to quickly verify the alignment of the latest agent ratings with human responses. **Data pollution.** Since our studies predate GPT-4V's latest training, data pollution was anticipated and became evident in our test on Imputation for Uncertainty using text and general descriptions. However, the agents' poor performance with specific image inputs made this issue less pronounced. Interestingly, data pollution offers an expected benefit. Though it hinders simulating new expert research, it aligns agent feedback with general evaluators for basic tasks, potentially making them useful for everyday visualization design and evaluation.

Knowledge-assisted and reliable agent simulation. As illustrated in our engineering efforts, injecting external knowledge can enhance the alignment of agent-based feedback in certain scenarios but raises concerns about generalizability and reliability. Content searched from websites could introduce additional biases, especially for non-expert evaluators unable to identify them. The challenge stems from the scattered and unsystematic nature of existing knowledge on human perception, with no comprehensive visual cognitive knowledge available. We attempted to use "Visualization Psychology" [20] and its references as a knowledge source for RAG to extract relevant information automatically. However, this failed due to semantic mismatches between visualization designs and knowledge, as well as uncertainty about relevant knowledge availability and selection. Reliable agent feedback will require advances in visualization research to provide systematic, practical knowledge for evaluation.

Profile-aware agents for diversity simulation. Agents reproduce patterns learned from large, general-purpose datasets, often reflecting the most common behaviors-what we term an "average user." Simply modifying demographic data, such as age or gender, does not affect agent outputs since such data is not emphasized during models' training. Due to privacy concerns, research often excludes detailed user profiles. To enable models to handle diverse user profiles, comprehensive first-hand data collection is necessary. For example, gathering complete user study logs from college students and fine-tuning large models could allow agents to emulate different student profiles. This method could be expanded to other groups, enabling profile-aware diversity in outputs and modeling intra- and inter-group variations.

Effective input preprocessing and prompting strategies. Our Study III shows that factors like modifying prompt keywords and mentioning scenes significantly influence agent ratings. Similar points have also been found in recent or concurrent work [3], [5], [14] when exploring models' accuracy and chart takeaways. While we did not systematically explore prompt optimization and all factor combinations, future work should carefully consider preprocessing and prompt strategies. Furthermore, because other open-source LMMs are less capable, making them unsuitable for thorough testing, we could not confirm the generalizability of our results. Future use of large modelbased agents will require preliminary studies to gather targeted insights for subsequent optimization.

Flexible agent-based procedure design. The design of our agent-based study was constrained by the current capabilities of agents, resulting in withinsubject and local comparisons. However, the emergence of advanced agents with visual capabilities and memory retention across multiple dialogue rounds could expand design possibilities. These agents could simulate a more human-like process by using memory to mimic how people rate images in a temporally continuous, between-subjects manner. Additionally, with scaling laws driving rapid improvements in large models and increasing context token capacities, future agents may overcome current limitations on the number of image inputs, enabling larger-scale withinsubject study designs. Still, inputting images close to the model's context limit—as in our Study II—may affect reasoning quality due to degraded long-range reasoning, and a tendency to overlook information in the middle of the input. Future procedure designs should also explore how context length usage impacts simulation outcomes.

Simulating broader feedback in broader scenarios. We investigated agent-based study procedures for quantitative data, applicable to broad evaluations using Likert-scale rating. While most OSF records lack user think-aloud data, some studies, like those on timelines [8], have documented participant strategies in natural language, warranting further investigation. In addition, our study excludes interactive or complex decision-making tasks and leaves a gap in evaluating interactive systems. For such systems, agents can be prompted to understand interfaces and tasks, thus enabling interaction sequences for task completion. This approach can be integrated with computational models to assess time overhead or combined with our study to simulate ratings. Systems requiring domain expertise and complex decision-making, such as visual analytics, demand agents equipped with domain-specific knowledge and advanced prompt chains. Future studies involving more advanced agents should focus on these complex tasks to enhance the depth of analysis.

Limitation and future work. Our study is limited to six existing studies due to the restricted availability of open-source resources. Additionally, although our incremental experiments complemented each other, they did not offer a complete picture. Specifically, we did not delve deeply into qualitative analysis of textual content [19], which could provide richer insights. In future work, we will conduct a systematic qualitative analysis of the textual data generated in our rating experiments to better understand the factors driving human-agent rating alignment. We also advocate for future empirical studies to include detailed experimental records to facilitate research on automated agent-based methods. Besides, the number of experts involved in the confidence coding stage was limited to five; future work should include a broader sample or cross-validation to strengthen the conclusions. Furthermore, our experiments are primarily limited to OpenAI's GPT series, and the conclusions cannot generalize to a wider range of models, training data, or architectures. We will utilize larger datasets and advanced techniques to keep pace with rapid advancements in the field.

Conclusion and Takeways

Based on all the findings, we answer the 3 RQs as:

- RQ1 (What are the characteristics of agent ratings): Agents can simulate ratings relatively aligned with human responses. Large modelbased agents are now aggregations of experiences, relying on a general understanding of visual graphs to mimic human outputs. Their ratings show promise in aligning with human ones, though they cannot yet capture profile diversity.
- RQ2 (How well do agent and human ratings align): The alignment of agent ratings is positively related to the confidence level of human experts for giving hypotheses. Agent ratings are more likely to align, only when our consulted external expert evaluators can offer consistent hypotheses with high confidence for the experiment.
- RQ3 (How to better align agent and human ratings): Constructive strategies applied to the agent inputs can influence alignment. Beyond directly utilizing materials in user study as the agent's textual and visual inputs, preprocessing them or integrating external knowledge could change rating alignment. Still, more research is needed for validation.

Given the insights and the potential scenarios we have studied, we identify three more key takeaways:

- **Potential Scenario.** Using large model-based agents in facilitating fast, iterative design of experiments is promising when their ratings can align with those of prototype experiments.
- Be Cautious. According to current results of current techniques of large models, agents cannot process

visual details like humans through vision or simulate diverse users. They cannot simulate most conclusions in visualization research since experts often lack confidence in their hypotheses and agent ratings are less likely to align with human data. We also need to cautiously handle adjustments to input and pipeline, mindful of their impact on alignment. Regardless, agent ratings require validation with real user data.

• **Promising Future.** Leveraging their strong reasoning, future advancements in large models' visual and cognitive functions promise wider applications for more aligned ratings and feedback across more scenarios.

Acknowledgments

We thank the anonymous reviewers for their valuable comments. This work is supported by the Natural Science Foundation of China (NSFC No.62472099) and Ji Hua Laboratory S&T Program (X250881UG250).

REFERENCES

- Liu, Shusen, Miao, Haichao, Li, Zhimin, Olson, Matthew, Pascucci, Valerio, Bremer, Peer-Timo, "AVA: Towards Autonomous Visualization Agents through Visual Perception-Driven Decision-Making," *Computer Graphics Forum*, vol. 43, no. 3, 2024.
- Yang, Fumeng, Ma, Yuxin, Harrison, Lane, Tompkin, James, Laidlaw, David H., "How Can Deep Neural Networks Aid Visualization Perception Research? Three Studies on Correlation Judgments in Scatterplots," *Proceedings of the CHI Conference on Human Factors in Computing Systems*, no. 822, pp. 1–17, 2023.
- Li, Zhimin, Miao, Haichao, Pascucci, Valerio, Liu, Shusen, "Visualization Literacy of Multimodal Large Language Models: A Comparative Study," *arXiv* preprint arXiv:2407.10996, 2024.
- Fan, Lizhou, Hua, Wenyue, Li, Xiang, Zhu, Kaijie, Jin, Mingyu, Li, Lingyao, Ling, Haoyang, Chi, Jinkui, Wang, Jindong, Ma, Xin, Zhang, Yongfeng. "Nphardeval4v: A dynamic reasoning benchmark of multimodal large language models," *arXiv preprint arXiv:2403.01777*, 2024.
- Wang, Huichen Will, Hoffswell, Jane, Thane, Sao Myat Thazin, Bursztyn, Victor S., Bearfield, Cindy Xiong, "How Aligned are Human Chart Takeaways and LLM Predictions? A Case Study on Bar Charts with Varying Layouts," *IEEE Transactions on Visualization and Computer Graphics*, vol. 31, no. 1, pp. 536–546, 2024.

- Adnan, Muhammad, Just, Mike, Baillie, Lynne, "Investigating Time Series Visualisations to Improve the User Experience," *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pp. 5444–5455, 2016.
- 7. OpenAI, "GPT-4 Technical Report," 2023.
- Di Bartolomeo, Sara, Pandey, Aditeya, Leventidis, Aristotelis, Saffo, David, Syeda, Uzma Haque, Carstensdottir, Elin, Seif El-Nasr, Magy, Borkin, Michelle A., Dunne, Cody, "Evaluating the Effect of Timeline Shape on Visualization Task Performance," Proceedings of the CHI Conference on Human Factors in Computing Systems, pp. 1–12, 2020.
- He, Tingying, Zhong, Yuanyang, Isenberg, Petra, Isenberg, Tobias, "Design Characterization for Blackand-White Textures in Visualization," *IEEE Transactions on Visualization and Computer Graphics*, vol. 30, no. 1, pp. 1019–1029, 2024.
- Reimann, Daniel, Blech, Christine, Ram, Nilam, Gaschler, Robert, "Visual Model Fit Estimation in Scatterplots: Influence of Amount and Decentering of Noise," *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 9, pp. 3834–3838, 2021.
- Sarma, Abhraneel, Guo, Shunan, Hoffswell, Jane, Rossi, Ryan, Du, Fan, Koh, Eunyee, Kay, Matthew, "Evaluating the Use of Uncertainty Visualisations for Imputations of Data Missing At Random in Scatterplots," *IEEE Transactions on Visualization and Computer Graphics*, vol. 29, no. 1, pp. 602–612, 2023.
- Bradley, Duncan, Strain, Gabriel, Jay, Caroline, Stewart, Andrew J., "Magnitude Judgements Are Influenced by the Relative Positions of Data Points Within Axis Limits," *IEEE Transactions on Visualization and Computer Graphics*, vol. 31, no. 2, pp. 1414–1421, 2025.
- Xu, Zhongzheng, Wall, Emily, "Exploring the Capability of LLMs in Performing Low-Level Visual Analytic Tasks on SVG Data Visualizations," *arXiv preprint arXiv:2404.19097*, 2024.
- Lo, Leo Yu-Ho, Qu, Huamin, "How Good (Or Bad) Are LLMs at Detecting Misleading Visualizations?," *IEEE Transactions on Visualization and Computer Graphics*, vol. 31, no. 1, pp. 1116-1125, 2025.
- Wang, Huichen Will, Gordon, Mitchell, Battle, Leilani, Heer, Jeffrey, "DracoGPT: Extracting Visualization Design Preferences from Large Language Models," *IEEE Transactions on Visualization and Computer Graphics*, vol. 31, no. 1, pp. 710-720, 2025.
- Cui, Zhenxing, Chen, Lu, Wang, Yunhai, Haehn, Daniel, Wang, Yong, Pfister, Hanspeter, "Generalization of CNNs on Relational Reasoning With Bar

Charts," *IEEE Transactions on Visualization and Computer Graphics*, pp. 1-15, 2024.

- Haehn, Daniel, Tompkin, James, Pfister, Hanspeter, "Evaluating 'Graphical Perception' with CNNs," *IEEE Transactions on Visualization and Computer Graphics*, vol. 25, no. 1, pp. 641–650, 2019.
- Cui, Yuan, Ge, Lily W., Ding, Yiren, Harrison, Lane, Yang, Fumeng, Kay, Matthew, "Promises and Pitfalls: Using Large Language Models to Generate Visualization Items," *IEEE Transactions on Visualization and Computer Graphics*, vol. 31, no. 1, pp. 1094-1104, 2025.
- Hämäläinen, Perttu, Tavast, Mikke, Kunnari, Anton, "Evaluating Large Language Models in Generating Synthetic HCI Research Data: a Case Study," *Proceedings of the CHI Conference on Human Factors in Computing Systems*, no. 433, pp. 1–19, 2023.
- 20. Szafir, Danielle Albers, Borgo, Rita, Chen, Min, Edwards, Darren J., Fisher, Brian, Padilla, Lace, "Visualization Psychology," Springer Nature, 2023.

Zekai Shao is a Ph.D. candidate at the School of Data Science, Fudan University, China. Contact him at zkshao23@m.fudan.edu.cn.

Yi Shan is a Ph.D. student at the School of Data Science, Fudan University, China. Contact him at ydan24@m.fudan.edu.cn.

Yixuan He is a master's student at the School of Data Science, Fudan University, China. Contact him at yixuanhe24@m.fudan.edu.cn.

Yuxuan Yao is an undergraduate student at the School of Data Science, Fudan University, China. Contact him at Yao1282557919@gmail.com.

Junhong Wang receives her Master's degree from the School of Data Science, Fudan University, China. Contact her at 20210980030@fudan.edu.cn.

Xiaolong (Luke) Zhang is an Associate Professor at the Pennsylvania State University, United States. Contact him at xuz14@psu.edu.

Yu Zhang receives his Ph.D. from the Department of Computer Science, University of Oxford, United Kingdom. Contact him at yuzhang94@outlook.com.

Siming Chen is an Associate Professor at the School of Data Science, Fudan University, China. Contact him at simingchen@fudan.edu.cn.