Tell Me Without Telling Me:
Two-Way Prediction of Visualization Literacy and Visual Attention

Minsuk Chang (iD), Yao Wang (iD), Huichen Will Wang (iD), Yuanhong Zhou,
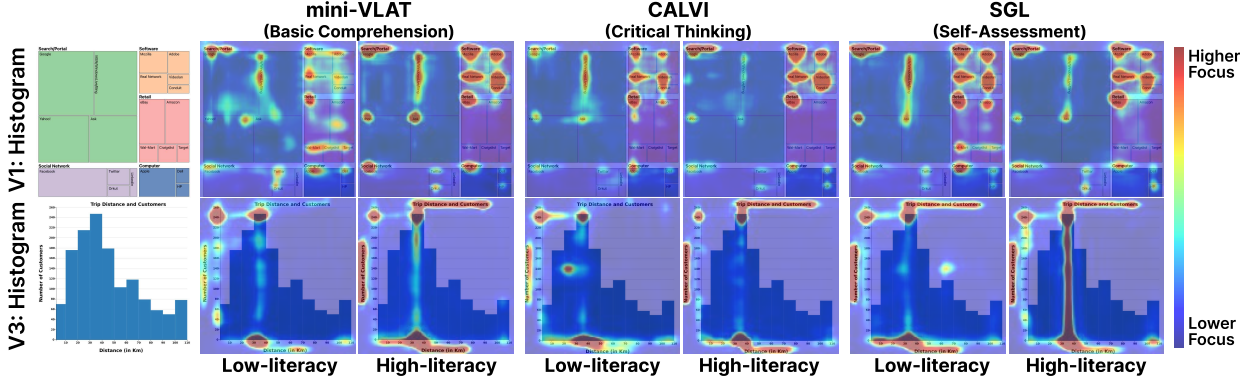Andreas Bulling (iD), and Cindy Xiong Bearfield (iD)

Fig. 1: Predicted visual saliency map from Lit2Sal model on three literacy assessment test (mini-VLAT, CALVI, SGL) scores for treemap and histogram. High-literacy group represents the top 20% on each test, and Low-literacy group represents the bottom 20%. The saliency map for high scores on mini-VLAT and CALVI focuses more on the title and axis labels than low scores, also showing more focused attention. Refer to Section 5 for further details.

Abstract—Accounting for individual differences can improve the effectiveness of visualization design. While the role of visual attention in visualization interpretation is well recognized, existing work often overlooks how this behavior varies based on visual literacy levels. Based on data from a 235-participant user study covering three visualization tests (mini-VLAT, CALVI, and SGL), we show that distinct attention patterns in visual data exploration can correlate with participants' literacy levels: While experts (high-scorers) generally show a strong attentional focus, novices (low-scorers) focus less and explore more. We then propose two computational models leveraging these insights: *Lit2Sal* – a novel visual saliency model that predicts observer attention given their visualization literacy level, and *Sal2Lit* – a model to predict visual literacy from human visual attention data. Our quantitative and qualitative evaluation demonstrates that Lit2Sal outperforms state-of-the-art saliency models with literacy-aware considerations. Sal2Lit predicts literacy with 86% accuracy using a single attention map, providing a time-efficient supplement to literacy assessment that only takes less than a minute. Taken together, our unique approach to consider individual differences in salience models and visual attention in literacy assessments paves the way for new directions in personalized visual data communication to enhance understanding.

Index Terms—Visualization Literacy Assessment, Visual Attention and Saliency, Visual Saliency Models.

---

1 INTRODUCTION

Visualization can guide users to notice key patterns in data. Yet what counts as the 'right' design often depends on who is looking. People interpret data on a deeply personal level [64,78], thus visualization design effectiveness can depend on their analytic tasks [48], their goals [67], and most importantly, their visualization literacy [58,62].

Existing research has demonstrated that we can improve visualization design by understanding where people look in a visualization [16]. Researchers have built saliency models to predict which parts of a visualization are most likely to attract a viewer's attention [15,18,53,76,85]. These models have proven valuable in supporting visualization and tool design, such as attention-aware UI [55], chart compression [16], and image quality evaluations [80].

However, existing models still largely assume universal viewing patterns among people and overlook differences driven by cognitive abilities such as visualization literacy [56,79]. We posit that **individuals with varying levels of visualization literacy interpret visualizations through distinct viewing patterns and can be effectively captured by models of their visual attention.** Therefore, we argue that we can improve existing saliency models by accounting for the unique patterns in viewers' visual attention depending on their literacy levels.

In this paper, we conduct user study (N = 235) using three established literacy tests: the mini Visualization Literacy Assessment Test (mini-VLAT) [62], the Critical thinking Assessment for Literacy in VIsualizations (CALVI) [35], and the Subjective Graph Literacy assessment (SGL) [34]. We recorded participants' responses and attention maps generated using the established BubbleView [40] technique, which approximates gaze tracking with mouse clicks. **Our analyses reveal striking attention differences between high- and low-literacy groups in each assessment**: Experts focus on specific regions of visualizations, creating concentrated 'hot spots' that reflect their targeted attention. In contrast, the viewing patterns of novices tended to be more distributed, with less intense focal points. This finding provides strong evidence that performance on visualization literacy tests does

- *Minsuk Chang, Yuanhong Zhou, and Cindy Xiong Bearfield is with Georgia Tech. E-mail: {minsuk, yzhou842, cxiong}@gatech.edu*
- *Yao Wang and Andreas Bulling is with University of Stuttgart. E-mail: {yao.wang, andreas.bulling}@vis.uni-stuttgart.de*
- *Huichen Will Wang is with University of Washington. E-mail: wwill@cs.washington.edu*

correlate with distinct visual attention patterns when interpreting visualizations. Informed by these insights, we then introduce two novel saliency models for literacy-aware attention prediction (**Lit2Sal**) and visual attention-based visualization literacy prediction (**Sal2Lit**).

Lit2Sal extends the state-of-the-art saliency model VisSal-Former [86] to predict where people look in a visualization while accounting for their literacy levels, across three literacy assessment categories: mini-VLAT [62] for visualization comprehension, CALVI [35] for critical thinking abilities, and SGL for self-perceived literacy, both individually and holistically. By integrating these measures, our model generates distinct saliency predictions for novices and experts.

Sal2Lit can predict visualization literacy levels for all three tests with an average accuracy of 86% with a human attention map from only one visualization. If we expand the input data to include attention maps from three visualizations, our model accuracy increases by over 87% (mini-VLAT) to 90% (CALVI, SGL). This time-efficient proxy supplements existing literacy assessments and reveals visual processes that differentiate novices and experts in visual data exploration.

Contributions: We conducted a crowdsourced study and a series of analyses, discovering how different aspects of visualization literacy (e.g., basic comprehension, critical thinking, and self-assessed ability) correlate with viewers' gaze patterns. Our literacy-to-saliency model, Lit2Sal, enhances existing saliency models by incorporating individual differences and aligning attention predictions with a viewer's visualization literacy level. Conversely, our saliency-to-literacy model, Sal2Lit, takes a perception-driven approach to reframe how visualization literacy can be assessed through observed visual attention patterns. Together, they advance our fundamental understanding of visual patterns in visual data exploration. By embracing personalization, these models empower researchers and practitioners to design personalized visualizations and efficiently estimate literacy levels.

2 RELATED WORK

2.1 Visual Attention in Visualizations

Where do people look in a visualization? Cognitive and vision scientists have long studied human attention to understand how visual features such as color, contrast, and orientation can guide attention patterns [39, 81]. With the rise of large-scale eye-tracking datasets and advances in machine learning, researchers have developed more sophisticated models to predict human attention patterns across images, text, and data visualizations. This research identifies salient regions that most likely attract human attention [6, 23, 74]. Saliency models serve as critical tools for evaluating the visual hierarchy in attention allocation in scene interpretation [39] and object identification [7].

While conventional saliency models can efficiently predict human attention patterns in natural scenes, they fail to account for the specialized perceptual and cognitive principles that govern our processing of data visualizations [25, 33, 41, 66]. This results in inaccurate predictions of attention patterns with visualized data. To address this, visualization researchers have developed saliency prediction models tailored to visualizations [18, 53]. However, these early attempts primarily relied on manually crafted rules based on pixel values, which can lose generalizability when the visualization gets complex and diverse. Recent advancements leverage deep learning techniques and extensive eye-tracking datasets in both natural images [50] and visualizations [32, 76] to improve prediction accuracy. Furthermore, innovative eye-tracking alternatives such as webcam-based tools [63] and mouse interactions [40, 59, 72] have enabled researchers to collect large-scale saliency data, overcoming the practical limitations associated with conventional eye-tracking approaches. Advances in visual attention modeling also considered factors such as task relevance [86] and gaze path sequences [84], leading to more nuanced models. These approaches typically aggregate saliency data from all participants to define a general consensus-based ground truth. We supplement these approaches to further improve visualization saliency models by considering individual data literacy levels in visual data interpretation,

Alongside the model development, various datasets have been collected to explore and analyze viewer saliency patterns in visualizations

more systematically [12, 44, 76]. Borkin et al. [8] collected an eye-tracking dataset with 393 visualizations from the MASSVIS [9] dataset. Shin et al. [76] crowdsourced saliency maps over 10,000 visualizations via webcam-based eye tracking. However, both datasets were collected in a task-agnostic setting, where saliency data was collected without considering the top-down influence on visual attention. To supplement the lack of task-specific saliency data, Gomez et al. [36] created a dataset of 20 visualizations with task-driven saliency, covering analytic tasks, including value retrieval in bar charts and outlier point detection in scatterplots. Polatsek et al. [66] analyzed how low-level tasks affect saliency through an eye-tracking study with 30 visualizations. Most recently, with a mouse-based proxy BubbleView [40], Wang et al. [86] collected the first large-scale saliency dataset containing task-relevance regions using 3,000 visualizations from ChartQA [52] dataset. However, no existing dataset has linked saliency patterns with visualization literacy, despite evidence that viewing patterns can predict cognitive abilities [22]. We address this gap by collecting saliency maps along with three literacy assessments, annotating literacy test performance from VLAT [62], CALVI [35], and SGL [34] to provide insights into the effect of individual differences in saliency modeling.

2.2 Visualization Literacy Assessments

Visualization literacy has long been a critical area of interest in the visualization community [11]. Early assessments of visualization literacy often relied on unstructured methods, primarily designed for specific domains, which involved manually curated visualization-question pairs [10, 45, 51, 71, 88]. For instance, Schönborn et al. [73] analyzed biochemistry students' ability to understand external representations and offer insights into domain-specific literacy. Baker et al. [2] and Huron et al. [38] measured the visualization literacy of participants by making them create or update the charts. Rodrigues et al. [70] explored how novices attempt to interpret visualizations by letting them ask questions rather than responding to predefined ones.

However, these domain-specific and exploratory methodologies limited the generalizability of research findings and comparability across different studies. To address this gap, Lee et al. [46] proposed VLAT, a standardized test assessing basic chart comprehension across diverse visualization types and tasks. Garcia et al. [34] created a standardized self-assessment inventory of visualization literacy based on chart familiarity. In addition to evaluating chart comprehension and familiarity, researchers also recognized that literacy could be assessed through higher levels of thinking, such as critical reflection of visualization takeaways and evaluations of argument quality from visualizations [13]. For example, to evaluate the critical thinking process of visualization literacy, Ge et al. [35] created CALVI by adding trick questions where misinformation is deliberately included. Quadri et al. [67] assessed the viewer's high-level chart comprehension by qualitatively comparing the takeaways with theme analysis. Adar and Lee [1] introduced a learning objectives framework to evaluate how viewers comprehend communicative visualizations matching the designer's goal. These assessments were frequently adopted in other research to assess general visualization literacy [29, 47], understand barriers better [21, 58], or benchmark large language models' visualization literacy [4, 37, 61].

Researchers also considered enhancing the reliability and efficiency of these standardized tests by slightly tuning their structures. For example, mini-VLAT, a subset of VLAT [46], was developed with comparable diagnostic capability. Cui et al. [26] also devised a compressed version of CALVI, which reduced the average test questionnaires by half [26]. Despite these efforts, assessing literacy remains time-intensive, particularly across multiple cognitive levels in Bloom's Taxonomy [31], from basic comprehension to critical thinking. This further motivated us to estimate literacy levels efficiently.

Although standardized tests provide individual score-based assessments, researchers typically group people to understand their common aspects better. For example, in the field of education, Persky et al. [65] proposed a framework comprising five stages of expertise, ranging from beginner to expert, drawing inspiration from the Dreyfus model [30], which similarly categorizes adults' skill acquisition into five distinct phases. Reding et al. [68] aimed to split the participants into four quan-
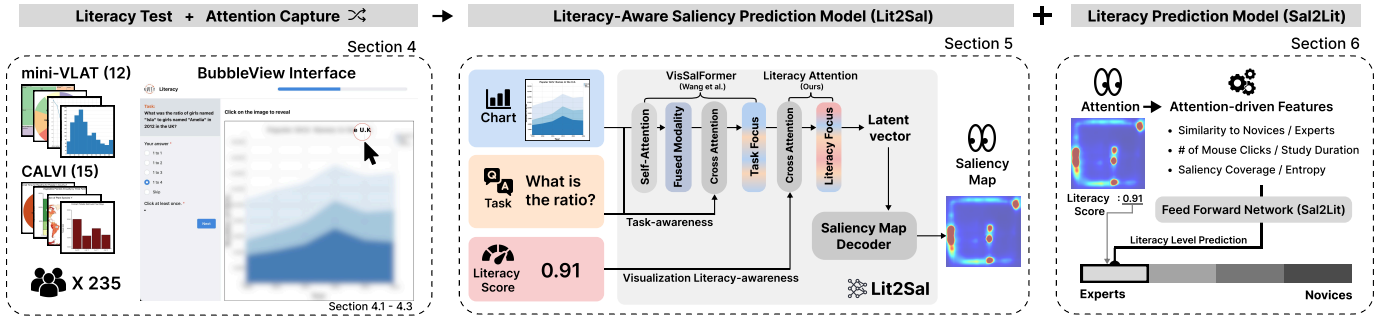
Fig. 2: Overview of our study from data collection to model experiments. We first collected attention data from 235 participants using BubbleView [40] while they solved 27 visualization literacy test questions consisting of mini-VLAT [62] and CALVI [35]. The data was used to test hypotheses in Section 4.4-4.6 and train two models: Lit2Sal (Section 5), and Sal2Lit (Section 6).

tiles but employed a binary classification of experts and novices due to the limited dataset size. For visualization literacy, proficiency has been self-assessed using three categories: novice, intermediate, and expert [82]. Building on this literature, we design a predictive model to assess literacy among two to five levels. Drawing from various definitions, we also use the terms experts and novices to refer to individuals with the highest and lowest visualization literacy, respectively, as measured by test scores.

3 HYPOTHESES

To understand how visualization literacy relates to visual attention, we formulate three hypotheses that examine this relationship in progressive stages. First, we discover how different aspects of literacy are measured and interrelated. Existing assessments of visualization literacy vary in the types of visualizations and questions they employ [34, 35, 46, 62], each capturing different aspects of literacy. We focus on three widely used assessments with different types of assessments: mini-VLAT [62] (basic comprehension), CALVI [35] (critical thinking), and SGL [34] (self-assessment), and hypothesize that

> **H1:** The distribution of participants' scores will differ across the assessments, reflecting the distinct aspects of literacy that each test is designed to measure.

Then, we analyze how attention pattern differs between groups of participants with different literacy levels. In learning science, experts and novices are known to exhibit different cognitive processes, particularly in how they retrieve and organize relevant information [56]. Given that visualization viewing patterns are task-dependent [19, 86] and can vary significantly based on individual characteristics such as verbal working memory [79], we hypothesize that

> **H2:** Visualization experts and novices will exhibit distinct viewing patterns, leading to differences in their generated attention maps.

Considering both **H1** and **H2** with an assumption that we could build a bidirectional mapping between visualization literacy level and visual saliency, we finally hypothesize

> **H3:** Visualization literacy and viewing patterns will be correlated, such that high-literacy participants exhibit viewing behaviors similar to experts, while low-literacy participants resemble novices.

Finding evidence for these hypotheses would give us a clear view of how visualization literacy and viewing patterns correlate, providing a foundation for designing more adaptive and user-aware visualizations.

4 EXPERIMENT

We conducted a crowdsourcing study to evaluate our hypotheses. Figure 2 shows study results being used to build our literacy-to-saliency prediction model (Lit2Sal) and saliency-to-literacy prediction model (Sal2Lit). The study was deployed through Revisit [28] and Prolific.

4.1 Participants

Using the pilot study data with 40 participants, we conducted a power analysis to identify the smallest effect size across all features listed in Section 4.6, which had a Cohen's f of 0.057. Following the power analysis with G*Power, we recruited 235 participants to detect an effect with 95% power at $\alpha = 0.05$. The participants had a mean age of $M_{age} = 38.73$ years ($\sigma_{age} = 13.03$), with 89 identifying as female. Regarding education, the plurality (89) held a bachelor's degree or higher, followed by 67 with a high school diploma or equivalent and 66 with a master's degree. Regarding visualization experience, 83 participants reported occasionally reading data visualizations, while 59 did so frequently. Additionally, 43 occasionally created visualizations for work or hobbies, and 30 reported often creating data visualizations. All participants were compensated at a rate of $12 per hour.

4.2 Methodology

Visualization Literacy Tests: Since visualization literacy can be assessed in multiple ways (Section 2.2), we employed three complementary approaches for a more comprehensive measurement: mini-VLAT [46] (basic literacy), abridged CALVI [35] (critical thinking), and SGL [34] (self-assessment).

We used the mini-VLAT [62] to assess basic comprehension to ensure the study runs for a manageable length. For critical thinking, because the full-length CALVI [35] consists of 47 trick questions and requires over 45 minutes to complete, we selected a 15-item subset with the highest discriminability, following the authors' recommendation in subsequent work by Cui et al. [26]. This combination of mini-VLAT and abridged CALVI contains both questions with and without misinformation, which Ge et al. [35] have demonstrated to support participants' critical thinking process. SGL, being a standalone, short self-assessment, was administered in its entirety.

Capturing Viewing Patterns through Attention Proxy We also showed each visualization on mini-VLAT and abridged CALVI with the BubbleView [40] interface to collect the viewing pattern when taking the tests. While traditional eye-tracking hardware is effective, it can be cumbersome for collecting large-scale attention data. Recent research has validated mouse-based proxies as reliable alternatives for capturing attention from images [17, 19, 40, 57, 59, 63, 72]. Therefore, we used a mouse-based proxy, BubbleView [40], for data collection, especially suited for capturing intentional attention. The bubble size was initially set to 32px and then adjusted based on each participant's screen. This ensured consistent bubble sizes across different images and display resolutions, aligning with the size of human foveal vision [40, 86]. From the collected stream of click records, we generated a continuous 2D attention map by applying a Gaussian blur to the bubbles, using a kernel with a sigma of 19px, approximately the same as 1 degree of visual angle [15, 40]. Our pilot study (40 participants) showed that 95% of participants answered the literacy questions within 85.02 seconds per chart. Therefore, we displayed each visualization to a participant for 90 seconds, allowing them ample time to visually explore and answer the mini-VLAT and CALVI literacy questions. This time limit also served

mini-VLAT
- V1: Treemap V2: Line Chart
- V3: Histogram V4: Bar Chart
- V5: Area Chart
- V6: Pie Chart
- V7: Scatterplot
- V8: Bubble Chart
- V9: Stacked Area Chart
- V10: Choropleth map
- V11: Stacked Bar Chart
- V12: Stacked 100% Bar

CALVI
- T3: Scale Range T20: Scale Order
- T25: Scale Order T26: Scale Functions
- T30: Missing Data T14: Scale Directions
- T40: Overplotting T10: Scale Directions
- T35: Aggregation T43: Cherry Picking
- T42: Missing Normalization
- T37: Inappropriate Aggregation
- T47: Misleading Annotations
- T48: Concealed Uncertainty
- T49: Misleading Annotations

SGL
- S1: How good are you at working with bar charts?
- S2: How good are you at working with line plots?
- S3: How good are you at working with pies?
- S4: How good are you at inferring the size of a bar in a bar chart?
- S5: How good are you at determining the difference between 2 bars in a bar chart?
- S6: How good are you at projecting a future trend from a line chart?
- S7: Are graphs easier to understand than numbers?
- S8: How often do you find graphical information to be useful?
- S9: To what extent do you believe in the saying 'a picture is worth one thousand words'?
- S10: When reading articles, how helpful do you find graphs that are part of a story?
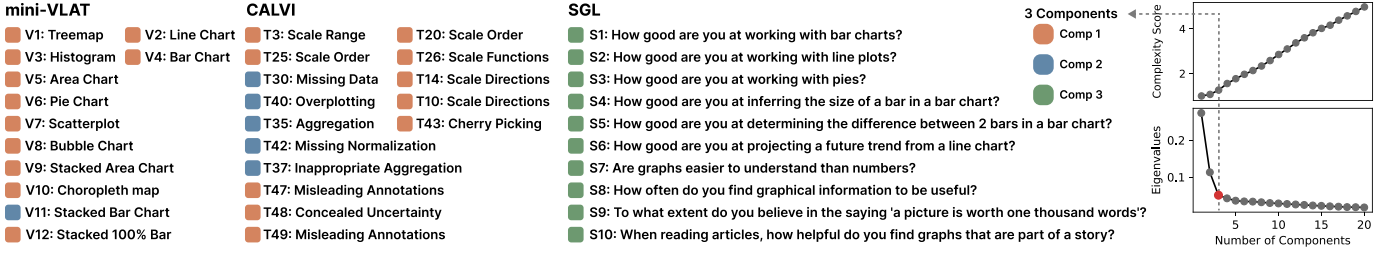
Fig. 3: Multiple Correspondence Analysis (MCA) results indicate three distinct components forming an elbow point in the component-eigenvalue chart. Each question is color-coded based on the component it contributes to the most.

as an attention check to encourage participants to efficiently engage with each question.

4.3 Procedure

The study was conducted as a between-subject experiment. After consenting to the study, participants self-assessed their visualization literacy with the SGL [34] test consisting of 10 questions with a 1-6 Likert scale. Then, we randomly presented them with 27 mixed multiple-choice questions from mini-VLAT and CALVI. The visualizations associated with each question were presented with a BubbleView interface, as shown in Figure 2 (left block). Each participant in main study had 90 seconds to freely explore the given visualization in BubbleView and answer the associated literacy question; pilot study had no time constraints. Also there were no specific instructions given on how to explore. After 90 seconds, they could only skip the question, but all interactions were logged. Participants could not return to previous questions. Upon completion, they provided demographic information: age, gender, education level, and experience in data visualization.

4.4 H1: Discriminability of Literacy Assessment Tests

We first analyzed how the three literacy assessment tests differ regarding overall score distribution and participant discriminability.

4.4.1 Overall Score Distribution

The score distributions with basic statistics are shown in Figure 4, where each score was normalized to the [0, 1] range based on its possible minimum and maximum values. For mini-VLAT, we first applied the score correction formula $R - W/(C - 1)$ (where R: correct, W: wrong, C: number of choices) to account for random guessing, as recommended by the authors [62], while CALVI and SGL used raw scores. Throughout the paper, we refer to the normalized scores for all three tests instead of the raw scores unless stated otherwise.

Each test score distribution resembled a normal distribution, with the mini-VLAT exhibiting the widest and most balanced spread. Regarding the balancedness of each test, we measured Pearson's third-moment coefficient of skewness. mini-VLAT and CALVI are slightly skewed to the left and right, respectively. SGL was highly left-skewed, showing a ceiling effect.

4.4.2 Relationships between Literacy Assessment Scores

The ceiling effect observed in SGL [34] correlates with performance on the mini-VLAT [62] with a resemblance to the Dunning-Kruger effect [20,42]. Figure 4 presents a regression plot of mini-VLAT scores against CALVI and SGL scores, highlighting this phenomenon. Participants with very low or moderately high mini-VLAT performance estimated their literacy to be low in the self-assessment (SGL). In contrast, those with slightly low or exceptionally high VLAT performance estimated their literacy to be high in SGL.

We formally characterize the relationship between mini-VLAT and SGL by analyzing the variance explained in a polynomial regression model. As the polynomial degree increases, the model captures more variance and becomes more complex. To balance model fit and complexity, we identify the elbow point where R^2 shows minimal improvement with additional degrees. This occurs at degree three, indicating that a third-degree polynomial provides the optimal trade-off.
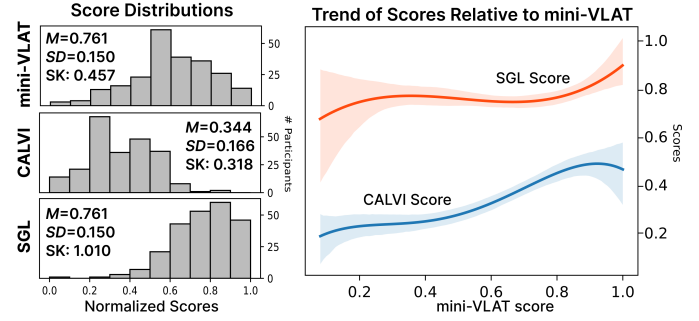

Fig. 4: (left) Score distribution for each test with statistics (average, standard deviation, and Pearson's 3rd skewness). All scores are normalized within [0, 1] based on the possible lowest and highest scores for each test. (right) Relationship between mini-VLAT score and other two scores (CALVI, SGL). The order of polynomial regression lines was calculated based on the elbow-point on the R^2 metric.

This allows us to express the relationship between SGL and VLAT performance as follows:

$$S = 0.598 + 1.216 \cdot V - 2.619 \cdot V^2 + 1.708 \cdot V^3 \tag{1}$$

where S refers to SGL and V to the mini-VLAT score. The extreme points where the Dunning-Kruger effect occurs are V = 0.36 and 0.67.

Similarly, we fit a polynomial regression model to describe the CALVI score with mini-VLAT. R^2 score had the elbow point on the fourth-degree polynomial, showing optimal trade-offs. The expressed relationship is:

$$C = 0.112 + 1.267 \cdot V - 4.852 \cdot V^2 + 7.989 \cdot V^3 - 4.046 \cdot V^4 \tag{2}$$

where C is the CALVI score. Analysis of the second derivative of the curve reveals a non-decreasing trend in C when the V ranges from 0 to 0.95. This suggests that, compared to SGL, CALVI shows a stronger positive correlation with mini-VLAT scores, indicating a relatively similar aspect of literacy measurement.

4.4.3 Multiple Correspondence Analysis (MCA)

We conducted a multiple correspondence analysis to examine further whether the three literacy tests capture distinct aspects of literacy, as they claim. For instance, mini-VLAT is focused on basic chart comprehension [62], CALVI is for critical thinking ability [35], and SGL is for self-assessment [34]. To uncover latent dimensions across the three tests, we applied MCA to item-level responses, using binary correctness for mini-VLAT and CALVI and a binarized form of SGL's six-level Likert scale. For SGL, we binarized the response by grouping the lower three levels as 'low' and the upper three as 'high'. As shown in Figure 3, the eigenvalues from the MCA indicate the variance explained by each additional component. At N = 3 components, we observe an 'elbow point' where each additional component after that no longer significantly explains more variances in the model despite increasing complexity. This motivates us to interpret this data with a three-factor model derived from MCA [69].

A closer examination of the component loadings from MCA reveals that all SGL items are positioned exclusively along Component 3,

Q: What is the average internet speed in Japan?   **Q: Does total number of employees in company Z decreased in 2022?**
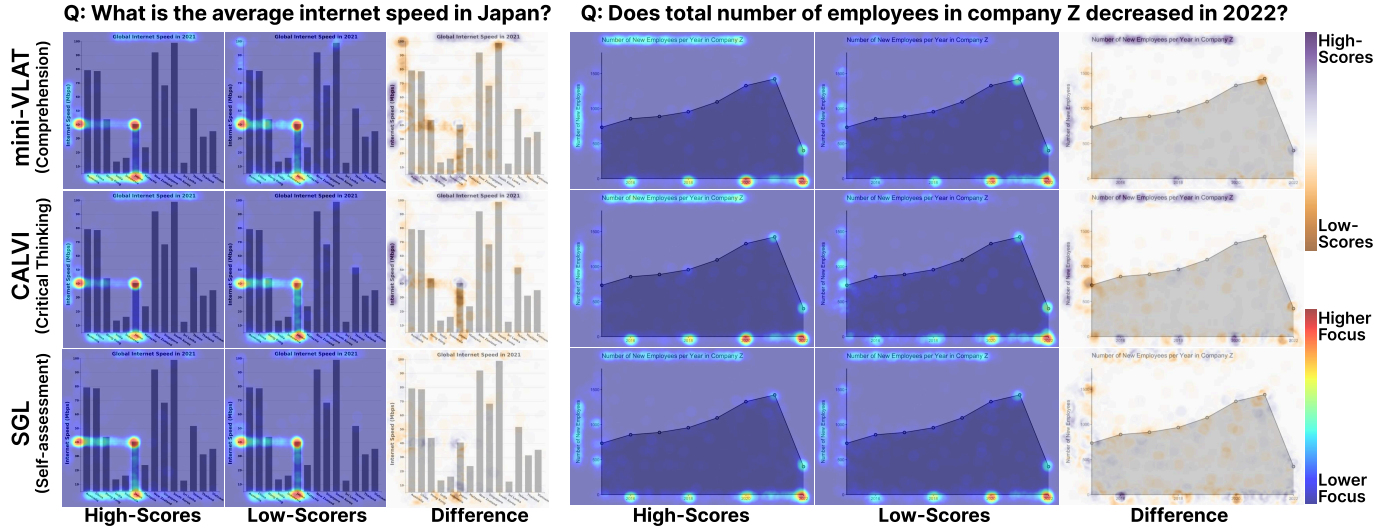

Fig. 5: Differences in visual attention between high- and low-scorers on three literacy tests (mini-VLAT, CALVI, and SGL) with V4 (Bar Chart) and T37 (Inappropriate Aggregation) charts. High and low scorers are defined as the top and bottom 25% of participants. In the attention maps, red regions indicate areas of higher visual attention, while blue regions indicate lower attention. The difference map highlights areas where high-performers exhibited greater attention in purple and where low-performers did in orange.
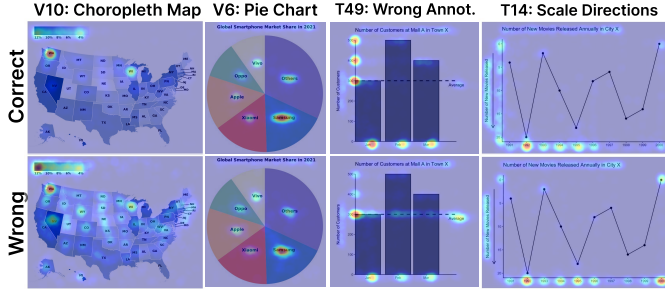

Fig. 6: Averaged attention map of participants who answered the task question correctly or not. The V10 and V6 were sampled from mini-VLAT, and T49 and T14 from CALVI. The chart codes are described in Figure 3.

Category	Group	Coverage	IoU$_{Pair}$	IoU$_{Title}$	IoU$_{Labels}$	IoU$_{Legend}$
Accuracy	Correct	**0.045**	**0.320**	**0.074**	**0.098**	**0.126**
	Incorrect	0.040	0.245	0.053	0.084	0.096
mini-VLAT	Experts	**0.045**	**0.299**	**0.090**	**0.128**	**0.124**
	Novices	0.039	0.239	0.042	0.060	0.094
CALVI	Experts	**0.046**	**0.294**	**0.092**	**0.119**	**0.123**
	Novices	0.038	0.243	0.040	0.057	0.094
SGL	Experts	0.041	**0.260**	0.055	0.087	0.102
	Novices	**0.045**	0.250	**0.068**	0.086	0.113

Table 1: Average metrics by accuracy and literacy groups. Coverage reflects the wideness of attention maps. IoU is computed for within-group pairs (IoU$_{Pair}$) or with specific chart elements (title, labels, legend). Bold indicates significance (p<0.001, pairwise t-tests). SGL's IoU with labels and legend was not significant (p>0.05).

with no overlap from mini-VLAT or CALVI items. This indicates that the self-perceived literacy measured by SGL represents a distinct construct, separate from the dimensions assessed by the other two tests. In contrast, mini-VLAT and CALVI items are distributed across Components 1 and 2, indicating the presence of two additional latent structures of visualization literacy. Each test appears to contribute to both components, implying that they capture overlapping but not identical aspects of comprehension and critical reasoning. Based on the assigned items, we interpret Component 1 as standard questions and Component 2 as difficult questions. Component 2 corresponds to a set of questions discriminating the expert level, as it includes five items from CALVI reported to be the most challenging among the 15 questions [35]. Among mini-VLAT questions, Component 2 includes the stacked bar chart with the lowest accuracy.

Takeaway: These analyses show that while the three literacy assessments are correlated, they each capture distinct aspects of visualization literacy, **supporting H1.**

4.5 H2: Viewing Patterns for Experts versus Novices

We adopt a two-step analysis to compare attention maps generated based on viewing patterns between visualization experts and novices. First, at the individual question level, we compare attention maps associated with **correct** versus **incorrect** responses to determine whether viewing patterns differ based on response accuracy. Second, we aggregate attention maps across multiple questions for each participant to examine whether overall high-scoring individuals (i.e., **experts**) exhibit systematic differences in viewing patterns compared to low-scoring individuals (i.e., **novices**) across all three literacy assessments. The

quantitative results are summarized in Table 1.

4.5.1 Difference of Attention by Question Accuracy

We first qualitatively analyzed how the visual attention maps differ based on the individual correctness of the visualization literacy test questions. Figure 6 shows that the **attention maps associated with correct answers were generated by participants who were more selectively focused**, attending to fewer regions but with high intensity. Additionally, participants who answered CALVI questions incorrectly were more likely to attend to misinformation factors, such as incorrect annotations or peaks. We then measured the spread of attention with Saliency Coverage [86] and within-group alignment with pairwise IoU [5]. Saliency Coverage quantifies how widely a participant's attention is distributed (e.g., percentage of nonzero pixels in their attention map), while IoU measures the spatial overlap between two maps (e.g., intersection over union of attended regions). **Correct group exhibited a slightly higher coverage and much higher pairwise IoU**, indicating greater consistency and overlap in their attention patterns. This pattern makes sense because BubbleView [40] reveals chart details through a blurred image, strategically identifying task-relevant areas while minimizing exploration of irrelevant regions offers a clear advantage under time constraints. In contrast, attention maps linked to incorrect answers tend to show low overlap between answers even though having a small coverage, indicating a higher variance of attention.

We also observed differences in the **types of regions** associated with high attention between maps generated by correct versus incorrect responses. After manually splitting the charts into segments containing a title, axis label, and legend, we calculate the IoU with attention maps from correct/incorrect groups. Attention maps from the correct answers

Feature Importance in Linear Regression & Sal2Lit

Legend: ■ Regression Coefficient (Linear Regression - Sec 4.6.1) ■ Feature Importance (Sal2Lit - Sec 6.3.2)

Within-Map Features

	mini-VLAT (Reg)	mini-VLAT (Sal2Lit)	CALVI (Reg)	CALVI (Sal2Lit)	SGL (Reg)	SGL (Sal2Lit)
Click Counts	0.26***	0.06	0.09**	0.05	0.17***	0.04
Task Duration	-0.09***	0.10	-0.08***	0.06	-0.05	0.05
Shannon Entropy (SE)	-0.03	0.11	0.04	0.07	-0.11***	0.05
Saliency Coverage (SC)	-0.02	0.08	-0.01	0.04	-0.01	0.04

Between-Group Features

	mini-VLAT Reg Novice	mini-VLAT Reg Expert	mini-VLAT Sal Novice	mini-VLAT Sal Expert	CALVI Reg Novice	CALVI Reg Expert	CALVI Sal Novice	CALVI Sal Expert	SGL Reg Novice	SGL Reg Expert	SGL Sal Novice	SGL Sal Expert
Mutual Information (MI)	-0.26***	0.18***	0.06	0.05	-0.31***	0.40***	0.04	0.04	-0.91***	0.78***	0.03	0.03
(neg) KL-Divergence (KL-D)	-1.51***	0.80***	0.08	0.08	-1.12***	0.64***	0.06	0.05	-0.82***	1.04***	0.06	0.03
Spearman's R Corr. (SRC)	0.09	0.50***	0.09	0.07	0.08**	0.51***	0.06	0.06	-0.63***	0.24**	0.04	0.03
Structural Similarity (SSIM)	-0.10*	0.05	0.07	0.05	-0.08	0.00	0.06	0.08	-0.06	0.12**	0.03	0.03
Mean Squared Error (MSE)	0.10	0.00	0.06	0.05	-0.06	0.08	0.04	0.06	0.05	-0.06	0.03	0.03

Fig. 7: Relative feature importance from linear regression (Section 4.6.1) and Sal2Lit model (Section 6.3.2). Darker color represents higher importance in both colors. The linear regression's red color represents the coefficients' absolute value with a corresponding significance level (*: p<0.05, **: p<0.01, ***: p<0.001), and blue represents the Sal2Lit model's feature importance calculated with Integrated Gradients [54].
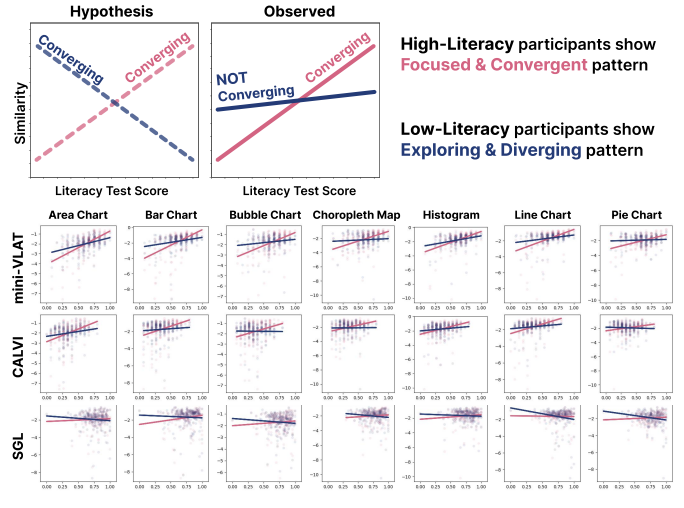

Fig. 8: (Top) Description of hypothesis on a similarity in attention map where experts and novices would have common areas in attention. (Bottom) The similarity of individual attention map against two groups calculated with negated KL-D. Only experts' attention converges into a certain pattern. See Section 4.6.2 for further analysis.

showed a greater focus on all three elements than those from the incorrect group, suggesting that accurate answers relied more on contextual information. We interpret these results as participants needing to remain aware that the current plot may contain misinformation from the mixed task orders from mini-VLAT and CALVI. Such awareness is also reflected in the y-axis of the bar and line charts in Figure 6.

4.5.2 Difference of Attention by Individual Literacy Scores

Moving beyond one-off question accuracy, we analyze participants' overall literacy scores to identify more generalizable attention trends across literacy levels. We sampled the top 25% and bottom 25% participants as experts and novices on three literacy tests, following Van et al. [68]'s quartile-based grouping of expertise.

As shown in Figure 5, there was a notable difference in primary focus areas across literacy levels. Similar to our question-level attention analysis comparing correct and incorrect responses, the expert group of all literacy tests had higher pairwise IoU than the novice group, showing highly overlapping focus. On the other hand, novice groups in all three tests had low pairwise IoU, showing a viewing pattern of spatial exploration. Also, the expert group of mini-VLAT and CALVI scores focused more on contextual elements (title, labels, legends), represented with higher IoU. In mini-VLAT and CALVI, novices also had low IoU against textual elements, showing more attention to graphical elements where data is embedded. However, SGL showed a different attention pattern, where novice participants focused more on the title than experts. Also, no significant difference in attention was found on labels and legends. This aligns with the findings in Section 4.4 that SGL measures a distinct dimension of literacy compared to mini-VLAT and CALVI.

Takeaway: Taken together, these findings **support H2**, suggesting that visual attention differs between novices and experts, driven by the distinct viewing patterns observed between correct and incorrect groups.

4.6 H3: Correlating Literacy and Viewing Patterns

To explore the relationship between visualization literacy and attention maps generated from viewing patterns, we described participant behaviors and quantified the attention maps via two complementary categories of predictive variables:

Within-Map Descriptive Features: We first adopt unitary features that describe the viewers' attention during the task. These are metrics extracted from individual attention maps, such as click count, task duration, Shannon entropy (SE), and saliency coverage (SC). They

capture a participant's attention distribution and focus within a single visualization, following methods from Wang et al. [86]. Higher SE and SC values indicate a more broad and dispersed attention map, reflecting less focused viewing patterns. Lower values indicate more concentrated attention.

Between-Group Comparative Features: Within-map features may only partially explain visual attention as they do not represent how the viewer's attention relates to certain groups' common behavior. Therefore, we also aim to evaluate how each participant's attention map aligns with others across different groups. We group the comparison maps by *literacy level* (expert vs. novice) and *task correctness* (correct vs. wrong). We apply five similarity metrics for each group to compare continuous attention maps.

They include: Mutual Information (MI) [83], Kullback-Leibler Divergence (KL-D) [43], Spearman's Rank Correlation (SRC) [77], Structural Similarity (SSIM) [87], and Mean Squared Error (MSE) [80] derived from Chang et al. [19]. These metrics capture different similarity aspects: MI and KL-D quantify information-theoretic divergence, SRC captures monotonic spatial correlation, SSIM assesses structural layout similarity, and MSE measures raw pixel-wise difference. For KL-D, we kept the group's averaged attention as the true probability distribution. By computing these similarity scores between each individual's maps and the aggregated version from high- and low-literacy groups, we can assess whether certain patterns are more diagnostic of literacy levels.

Together, these two categories offer a comprehensive and explainable mapping between attention patterns and literacy to test our hypothesis.

4.6.1 Feature Importance in Linear Regression

We performed a linear regression to predict literacy scores using these variables. The results are shown in Figure 7, highlighting the significance and strength of each feature, with all features normalized to a [0, 1] range. Each model for assessment scores partially showed power for explaining the variance in terms of R-squared (mini-VLAT: 0.397, CALVI: 0.341, SGL: 0.273).

On within-map features, click counts had a positive correlation, and task duration had a slightly negative correlation with literacy scores, both highly significant. Further analysis of the attention maps showed that experts clicked more frequently in task-relevant, high-focus regions. This underlines the finding that they focused selectively on fewer areas with greater intensity.

Among the similarity metrics on literacy level groups (experts and novices) of each literacy test, KL-D and MI had the strongest correlation, showing significance in all literacy types. Other metrics (SRC,

SSIM, MSE) had significance on some expertise groups and test scores, but the coefficients were not strong enough compared to the first two. Interestingly, only KL-D and MI negatively correlated with the novice group in all three test scores, showing how defining similarity matters in interpreting viewing patterns.

4.6.2 Task-Based Analysis of Attention Similarity and Literacy

Having shown that several features, including KL-D, significantly correlate with literacy scores, we further examine the relationship between literacy and viewing patterns by comparing each participant's attention map to group-level maps from experts and novices on a per-question basis. This analysis allows us to test whether participants with similar literacy scores may exhibit convergent viewing patterns. We conduct this analysis at the question level because, as shown in Section 4.4, some questions or charts may better differentiate between literacy levels, leading to more substantial convergence in viewing patterns among similarly scoring participants.

In Figure 8, we plot the relationship between literacy scores (x) and attention map similarity (y), which describes how strongly the viewing patterns converge to that of the experts (red) and novices (blue). We found that **high-scoring participants tended to show a high similarity of attention pattern with the expert group's map**, showing a positive slope (╱) in all three tests, supporting H3. However, **the attention pattern of low-scored participants did not appear similar to the novice group's**, shown as almost flat slopes (╌) in mini-VLAT and CALVI tests, providing evidence against H3. This further reinforces our findings that participants with high literacy tended to explore the visualization more selectively. In contrast, participants with low literacy tended to explore the image more widely and independently, so the aggregated pattern does not necessarily resemble the individual maps.

Conclusion: We discovered that experts and novices show distinct viewing patterns, quantified by a linear regression test that describes literacy scores with features derived from visual attention. These results now serve as a foundation for a predictive model that infers how someone views a visualization based on their literacy score.

5 LITERACY-AWARE SALIENCY PREDICTION

We trained a saliency prediction model, **Lit2Sal**, to predict saliency based on task and literacy levels to allow more personalized predictions.

5.1 Dataset Preparation

We trained our model with our collected saliency dataset of 235 participants on 27 chart-task pairs with three literacy score types (mini-VLAT, CALVI, SGL). To ensure a balanced training dataset, we applied oversampling to obtain an equal number of participants from each bin. Specifically, we divided the scores into five quantile-based bins, aligning with the maximum level of expertise differentiation identified in prior work [30, 65]. For evaluation, we formed a test split as 10% of the full dataset following Wang et al. [86]. We randomly sampled four participants from each quantile, resulting in 20 (participants) × 27 (charts) = 540 test data points, accounting for approximately 10% of the dataset.

5.2 Model Training

We created Lit2Sal's model structure by augmenting the original implementation of VisSalFormer [86]. We added an information-fusing component in the middle of the model structure to allow the conditional generation of a saliency map based on the literacy score.

5.2.1 Model Architecture

We mainly referred to VisSalFormer [86] as our baseline structure, the state-of-the-art visualization saliency prediction model. While keeping the Swin-transformer [49] as the image encoder and BERT [27] as the text encoder, we added an additional input of visualization literacy score as a single scalar value. As shown in the middle panel of Figure 2, the original fusion layer of VisSalFormer [86] merged visual (chart-based) and textual information (task-based) via the first two attention modules. This enables task-focused predictions that can predict viewer

Model	PCC ↑	NSS ↑	AUC ↑	SIM ↑	KL ↓
TSN (Pre-trained)	0.156	0.714	0.677	0.154	2.689
TSN (Literacy)	0.251	1.209	0.773	0.142	2.612
VSF (Pre-trained)	0.118	0.551	0.637	0.142	2.665
VSF (Literacy)	0.347	1.500	0.833	0.188	2.327
Lit2Sal (mini-VLAT)	0.429	<u>2.147</u>	<u>0.864</u>	0.255	<u>2.100</u>
Lit2Sal (CALVI)	**0.460**	**2.320**	**0.874**	0.254	2.170
Lit2Sal (SGL)	<u>0.430</u>	2.052	0.858	**0.257**	**2.086**

Table 2: Model performance comparison on five metrics with pre-trained weights or trained version with our dataset for TranSalNet (TSN) and VisSalFormer (VSF). All metrics for Lit2Sal with all literacy levels have significant differences (p < 0.001 except KL-D: 0.0016, AUC: 0.0022) compared to the VSF model (literacy-trained) based on pairwise t-tests.

attention based on their tasks. We appended a second cross-attention module that adds literacy scores to the prediction as a query vector. This allows the literacy score to additionally shape the latent vector.

5.2.2 Hyperparameters

We trained Lit2Sal on top of the VisSalFormer's default setting using the Adam optimizer with a learning rate of 5×10^{-5} for 300 epochs and a batch size of 128. The loss function was calculated with a combination of KL-divergence (KL-D), Normalized Scanpath Saliency (NSS), and Pearson's Correlation Coefficient (PCC), with a ratio of 10:5:2.

5.3 Evaluation

We evaluated Lit2Sal's performance against two state-of-the-art saliency prediction models capable of generating maps that align with human visual attention on images. The first is VisSalFormer [86], our backbone model for task-specific visual saliency prediction. The second is TranSalNet [50], an encoder-decoder Convolutional Neural Network specifically developed to predict saliency that closely follows human visual attention patterns.

5.3.1 Quantitative Results

We evaluated the generated saliency maps by measuring the KL-divergence (KL-D), Pearson's correlation coefficient (PCC), and histogram-based similarity (SIM) to assess their alignment with actual participants' visual attention maps. Because experts and novices produced different attention maps, we assessed their alignment with experts and novices depending on the literacy score input. Additionally, we used the Normalized Saliency Scanpath (NSS) and Area Under the Curve (AUC) to quantify how well the generated maps corresponded to human visual attention maps. We tested the TranSalNet and VisSalFormer models twice: once with their pre-trained weights and again after fine-tuning them on our literacy dataset. For both models, we averaged the saliency maps of participants with varying literacy scores in the training data, as they were designed to produce literacy-agnostic saliency predictions.

Table 2 presents the evaluation results, demonstrating that Lit2Sal achieves the highest performance across all metrics for all score types. This result provides strong evidence for literacy-aware saliency predictions. Additionally, we conducted a t-test for each metric to compare our model's performance with the second-best model, the fine-tuned VisSalFormer. Assuming Bonferroni adjustments for multiple-comparisons, the analysis revealed the **statistically significant superiority of the Lit2Sal model across all metrics** (p < 0.001 except KL-D: 0.0016; AUC: 0.0022). Note that our metric ranges differ from the reported numbers in the original model papers because we treated each participant's saliency as a single data point without aggregation, which led to more discrete maps and higher KL-D measures.

5.3.2 Qualitative Results

The predicted saliency maps from the test set are visualized in Figure 1, where expert and novice saliency maps were predicted with top 20% (first bin) and bottom 20% (last bin) ranking scores. Saliency predictions based on high literacy scores produced more focused maps with fewer noisy areas than those based on low literacy scores. Also, similar
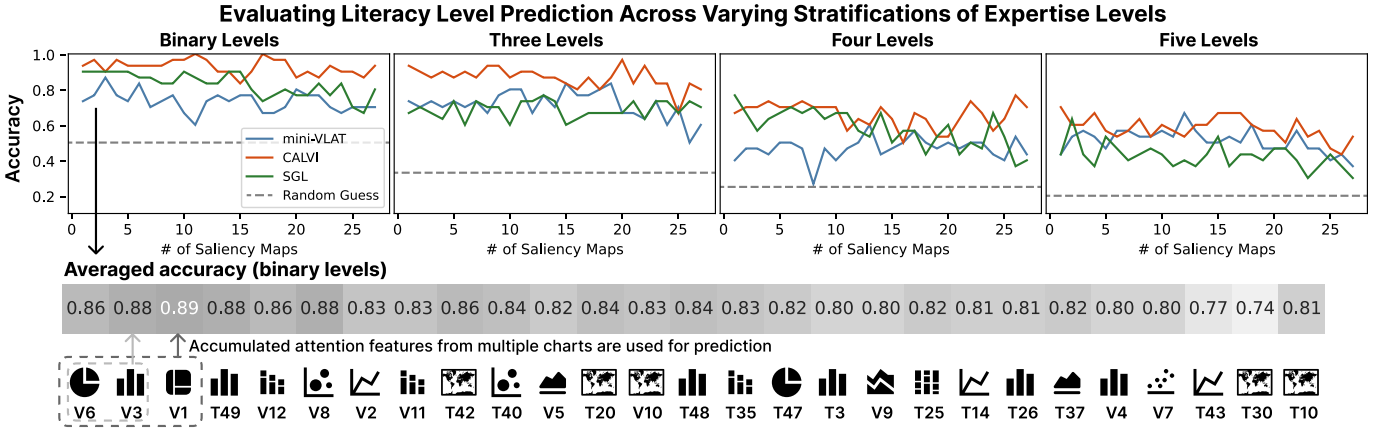
Fig. 9: Prediction results of participants' literacy levels using attention map-driven features from one or more charts. The four line charts illustrate prediction accuracy across different numbers of literacy levels, with baseline accuracy indicated by random guessing. The x-axis denotes the number of attention maps used for prediction. The heatmap below details the accumulative features from charts contributing to each performance level. Refer to Figure 3 for chart item codes.

to the pattern observed in Figure 5, we could find that high-literacy people in mini-VLAT and CALVI tend to focus more on the title and axis labels than novices. This could be interpreted as a strategy to solve tricky questions in the CALVI test, where the axis or legends often contain misinformation [35].

Conclusion: Our Lit2Sal model empirically maps continuous literacy scores to distinct saliency patterns, serving as a foundational step toward saliency prediction models that account for individual differences.

6 VISUALIZATION LITERACY PREDICTION

Following the successful training of the Lit2Sal model, we conducted another experiment to predict literacy levels using the attention maps, combining VLAT, CALVI, and SGL. We named this model **Sal2Lit**, where visual attention-driven saliency predicts the literacy level. The diagram of this model is shown in the right panel of Figure 2.

6.1 Dataset Preparation

Our dataset size for literacy level prediction is relatively small, consisting of only 235 participants' scores as individual data points. Therefore, we used all the within-map descriptive features (4) and the between-group comparative features (20) as described in Section 4.6 instead of raw attention maps to support the pattern discovery by our model.

Based on our preliminary experiment, predicting the exact score for all tests was difficult when only visual attention information was used. Inspired by use cases where expertise levels were analyzed by groups (see (Section 2.2)), we tested four versions of our Sal2Lit model, each predicting literacy scores as a categorical variable with two to five levels. The levels were determined based on score distribution quantiles. This allowed us to balance the number of data points in each quantile with oversampling, leading to more robust predictions. Also, following the evaluation setup of Sal2lit, we sampled the same 20 participants (four from each quantile) for the test set.

6.2 Model Training

Our Sal2Lit model was trained to predict three classes of literacy scores: mini-VLAT (basic comprehension), CALVI (critical thinking), and SGL (self-assessment). We also used a greedy algorithm [24] to identify the literacy test chart(s) that generate the most predictive visual attention maps, optimizing accuracy and minimizing data usage.

6.2.1 Model Architecture & Hyperparameters

Based on our preliminary experiments, traditional machine learning models such as Random Forest and Gradient Boosting proved insufficient for capturing the complexity of the problem. To address this, we employed a deep learning approach using a feedforward neural network with stacked linear layers, enabling the model to learn hierarchical feature representations.

Sal2Lit's network structure consists of five layers, with the hidden dimension starting at 512 and progressively halving until it reaches 32 before mapping to the final output layer corresponding to the number of classes (literacy levels). Our user study result showed that literacy test scores are correlated. Therefore, we aimed to predict all three literacy test levels simultaneously with a single model. To balance the prediction of the three levels, we averaged three individual cross-entropy loss functions with equal weights. We used the Adam optimizer with a learning rate of 1×10^{-4} for optimization. A batch size of 256 was chosen to maximize the GPU memory usage. The model was trained for a maximum of 150 epochs, with early stopping implemented to mitigate overfitting.

6.2.2 Greedy Subset Selection for Influential Charts

We aim for our visual attention-based model to supplement literacy tests by efficiently predicting literacy using only a few charts, minimizing time and effort. Exhaustively evaluating all possible chart item combinations is computationally expensive, so we applied a greedy selection algorithm to approximate the optimal subset. In this approach, we first identified the single chart item with the highest prediction accuracy when used alone. This was the bar chart from the mini-VLAT (V6), as shown in Figure 9. We then iteratively added one chart item at a time, each time selecting the item that produced the most significant improvement in prediction accuracy when combined with the previously selected items. Accuracy was measured as the average across the three literacy tests. This process continued until every chart item was added. Through this strategy, we efficiently identified a high-performing set of charts for literacy prediction without evaluating all possible combinations.

We also leave how to average three literacy scores' accuracy as a controllable parameter. For example, in some cases, the literacy level of comprehension (VLAT) would be preferred over critical thinking (CALVI), where all charts are proven to have no misinformation. Based on these scenarios, we can apply a weighted average to prioritize VLAT over CALVI and SGL.

6.3 Evaluation

We evaluated the trained Sal2Lit model based on the accuracy of its prediction results. The summarized results are plotted in Figure 9.

6.3.1 Overall Accuracy

Matching the intuition, the accuracy for binary classification of literacy level was generally higher than on more diverse levels in all three tests. Our model could predict the binary literacy level with **an average accuracy of 86% in all three literacy levels using only one visual attention map** (mini-VLAT: 73%, CALVI: 93%, SGL: 90%). Moreover, the accuracy exceeded 87% across all three tests using only three visualizations' attention maps: critical thinking (CALVI) reached 97%,

self-assessment (SGL) 90%, and basic comprehension (mini-VLAT) 87%. However, contrary to the expectation that incorporating information from additional charts would enhance performance, accuracy did not improve significantly with added charts and visual attention maps. This indicates a limit to the meaningful information that can be extracted from visual attention maps, beyond which the model begins to learn noise rather than meaningful patterns.

Comparing the results among the literacy types, we observe that CALVI's prediction accuracy is almost always higher than mini-VLAT and SGL. This suggests that the **difference in critical thinking ability is shown most vividly** as the difference in the attention pattern. In comparison, the mini-VLAT's accuracy was significantly lower than the other two tests, suggesting that the **general comprehension level is more difficult to capture** only based on the visual attention patterns.

The optimal charts and tasks selected by the greedy algorithm are shown in the bottom of Figure 9, which illustrates the ranking of the charts that contributed most to distinguishing experts from novices in literacy level prediction based on visual attention maps. The top three items (V6, V3, V1) with the highest average accuracy were from the mini-VLAT test. This suggests that the critical thinking ability measured by CALVI, can be captured even through VLAT questions without relying on trick questions. We discuss this deeper in Section 7.

6.3.2 Analyzing Feature Importance

To support our trained Sal2Lit model's robustness of its prediction performance, we analyzed which feature (from Section 4.6) contributed the most to the model's prediction. We applied the Integrated Gradient (IG) [54] technique to measure how each feature strongly affects the gradient flow. The impact is calculated by aggregating the gradient value in our model F, starting from the baseline feature $V_0 = 0$ to our test data's feature V. This is formulated as:

$$\mathbf{IG}(\mathbf{V}) = \mathbf{V} \int_0^1 \nabla_{\mathbf{V}} F(\alpha \mathbf{V}) \, d\alpha \qquad (3)$$

The result of applying IG to the binary-level prediction model is plotted in Figure 7 as a blue color scale of the predictive model. Among **the within-map descriptive features** (see Section 2.2), click counts had the highest regression coefficient (computed from our user study results), but Shannon entropy had the highest feature importance (computed from our Sal2Lit model) despite a low regression coefficient. This discrepancy between empirical user-driven results and model-based interpretation suggests a non-linear relationship of entropy value predicting literacy. Among **between-group comparative features**, KL-D showed both high feature importance and regression coefficients. Note that the calculated feature importance reflects the behavior of the trained model and does not necessarily indicate generalizable knowledge.

Conclusion: Our Sal2Lit model accurately predicts visualization literacy using features from as few as three visual attention maps, validating our hypotheses that viewing patterns can effectively distinguish between varying levels of literacy.

7 DISCUSSION

Several findings in our study offer new insights into how visual attention reflects literacy and critical thinking.

Implications for adaptive visualization and scaffolding. Our study reveals diverging behavior among literacy levels: experts show focused, 'convergent' attention on contextual elements like titles, while novices display dispersed, divergent' visual exploration. We reflect on our findings to consider their implications for adaptive visualization design.

While our current findings do not allow for any causal conclusions about whether novices become disoriented within the visualization or whether the visualization itself fails to guide them, we propose two plausible explanations. On one hand, novices may be more susceptible to cognitive overload when faced with unfamiliar visual structures, leaving them uncertain about where to begin extracting information. On the other hand, the visualization might not be designed to help certain viewers locate key patterns, increasing the likelihood of incorrect responses on comprehension or critical thinking questions. In both

scenarios, an adaptive system that detects disoriented viewing patterns could intervene with visual cues, like highlighting axes or annotations, to support user understanding. However, we acknowledge the risk that such designs may reduce user agency. As every visualization is a rhetorical device with multiple possible interpretations [3], adaptive systems must be designed with careful attention to the story the visualization intends to convey.

Redefining critical thinking in visualization. In our experiment, viewing patterns on mini-VLAT charts, which contained no misinformation, were strongly predictive of CALVI test scores. This finding suggests the presence of expert-like viewing behaviors associated with critical thinking, particularly those involving scrutiny of a chart's structure. These results point to the potential for a novel class of literacy assessments: rather than relying solely on standardized tests, researchers could evaluate critical thinking by analyzing viewing patterns on visualizations, offering a means to detect skepticism (and potentially even perceived credibility) through visual behavior.

8 LIMITATIONS & FUTURE WORK

We trained models to differentiate visualization novices from experts using visual attention maps from visual analytic tasks. However, our data collection method of BubbleView [40] is only a proxy of visual attention. BubbleView requires slow, deliberate clicking at a rate far slower than natural eye movements. This may disrupt natural visual exploration and increase cognitive load, reducing response accuracy [75]. Future work can test our models with eye-tracking data [63] and fine-tune the outcome to reveal a deeper relationship between attention patterns and literacy.

Further, visual literacy can be multi-faceted. Our current experiment tested only three tests: mini-VLAT, CALVI, and SGL. Future work can improve our seminal work by training the models on more extensive chart datasets with more participants that cover a broader spectrum of tasks and literacy levels. Such expanded evaluations would increase the performance and generalizability of our model across varied contexts

Finally, we found that similar visual attention patterns do not always indicate the same answers on a literacy test (see Figure 6 Pie Chart). Participants can focus on the same regions, producing both correct and incorrect responses. This suggests that attention does not guarantee comprehension, as misinterpreted prompts or time pressure can lead to errors [14,60]. This highlights the limitations of relying solely on visual attention as a proxy for cognitive processing and visual literacy. Future work should explore additional diagnostic cues that more accurately reflect literacy levels. Therefore, we claim that our **Sal2Lit model complements, rather than replaces prior literacy assessments.**

9 CONCLUSION

In this work, we conducted a user study investigating the relationship between visualization literacy and human visual attention patterns. Our results revealed distinct human visual attention patterns based on different literacy levels: experts selectively focused on fewer regions but with greater intensity, while novices explored more broadly. Building upon these findings, we introduced two computational models: **Lit2Sal**, a literacy-aware saliency model predicting where people will attend to in a visualization depending on their literacy levels, and **Sal2Lit**, which predicts one's visualization literacy level from their visual attention maps. Together, these models open opportunities for real-time adaptive visualizations that account for individual differences by customizing visuals to match viewers' literacy. Also, our approach has promising educational applications, such as attention-based feedback mechanisms that guide novices to attend to task-relevant regions. Finally, our models show potential for integration into design recommendation tools that dynamically evaluate and suggest improvements during the visualization authoring process, enhancing clarity and comprehension for audiences with diverse literacy levels.

REFERENCES

[1] E. Adar and E. Lee. Communicative visualizations as a learning problem. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):946–956, 2021. 2

[2] R. S. Baker, A. T. Corbett, and K. R. Koedinger. Toward a model of learning data representations. In *Proceedings of the 23rd annual conference of the Cognitive Science Society*, pp. 45–50. Mahwah NJ, 2001. 2

[3] C. X. Bearfield, L. Van Weelden, A. Waytz, and S. Franconeri. Same data, diverging perspectives: The power of visualizations to elicit competing interpretations. *IEEE Transactions on Visualization and Computer Graphics*, 2024. 9

[4] A. Bendeck and J. Stasko. An empirical evaluation of the GPT-4 multi-modal language model on visualization literacy tasks. *IEEE Transactions on Visualization and Computer Graphics*, 2024. 2

[5] A. Boggust, B. Hoover, A. Satyanarayan, and H. Strobelt. Shared interest: Measuring human-AI alignment to identify recurring patterns in model behavior. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pp. 1–17, 2022. 5

[6] C. E. Bonhage, J. L. Mueller, A. D. Friederici, and C. J. Fiebach. Combined eye tracking and fMRI reveals neural basis of linguistic predictions during sentence comprehension. *Cortex*, 68:33–47, 2015. 2

[7] A. Borji, M.-M. Cheng, H. Jiang, and J. Li. Salient object detection: A benchmark. *IEEE Transactions on Image Processing*, 24(12):5706–5722, 2015. 2

[8] M. A. Borkin, Z. Bylinskii, N. W. Kim, C. M. Bainbridge, C. S. Yeh, D. Borkin, H. Pfister, and A. Oliva. Beyond memorability: Visualization recognition and recall. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):519–528, 2015. 2

[9] M. A. Borkin, A. A. Vo, Z. Bylinskii, P. Isola, S. Sunkavalli, A. Oliva, and H. Pfister. What makes a visualization memorable? *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2306–2315, 2013. 2

[10] 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

[11] 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

[12] M. Burch, L. Chuang, B. Fisher, A. Schmidt, and D. Weiskopf. *Eye Tracking and Visualization: Foundations, Techniques, and Applications. ETVIS 2015*. Springer Publishing Company, Incorporated, 1st ed., 2018. 2

[13] A. Burns, C. Xiong, S. Franconeri, A. Cairo, and N. Mahyar. How to evaluate data visualizations across different levels of understanding. In *2020 IEEE Workshop on Evaluation and Beyond-Methodological Approaches to Visualization (BELIV)*, pp. 19–28. IEEE, 2020. 2

[14] R. F. Burton. Multiple-choice and true/false tests: myths and misapprehensions. *Assessment & Evaluation in Higher Education*, 30(1):65–72, 2005. 9

[15] Z. Bylinskii, T. Judd, A. Oliva, A. Torralba, and F. Durand. What do different evaluation metrics tell us about saliency models? *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(3):740–757, 2018. 1, 3

[16] Z. Bylinskii, N. W. Kim, P. O'Donovan, S. Alsheikh, S. Madan, H. Pfister, F. Durand, B. Russell, and A. Hertzmann. Learning visual importance for graphic designs and data visualizations. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*, pp. 57–69, 2017. 1

[17] Z. Bylinskii, A. Newman, B. McNamara, M. Tancik, S. Madan, F. Durand, and A. Oliva. ZoomMaps: Using zoom to capture areas of interest on images. In *Vision Sciences Society (VSS) Conference*, 2019. 3

[18] Z. Bylinskii, A. Recasens, A. Borji, A. Oliva, A. Torralba, and F. Durand. Where should saliency models look next? In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part V 14*, pp. 809–824. Springer, 2016. 1, 2

[19] M. Chang, Y. Wang, H. W. Wang, A. Bulling, and C. Xiong. Grid Labeling: Crowdsourcing task-specific importance from visualizations. In *Proceedings of the EuroVis 2025 Short Papers*. Eurographics Association, 2025. 3, 6

[20] M. Chen, Y. Liu, and E. Wall. Unmasking Dunning-Kruger effect in visual reasoning & judgment. *IEEE Transactions on Visualization and Computer Graphics*, 2024. 4

[21] K. Choe, C. Lee, S. Lee, J. Song, A. Cho, N. W. Kim, and J. Seo. Enhancing data literacy on-demand: Llms as guides for novices in chart interpretation. *IEEE Transactions on Visualization and Computer Graphics*, 2024. 2

[22] C. Conati, S. Lallé, M. A. Rahman, and D. Toker. Comparing and combining interaction data and eye-tracking data for the real-time prediction of user cognitive abilities in visualization tasks. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 10(2):1–41, 2020. 2

[23] K. Conklin and A. Pellicer-Sánchez. Using eye-tracking in applied linguistics and second language research. *Second Language Research*, 32(3):453–467, 2016. 2

[24] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, Cambridge, MA, 3rd ed., 2009. 8

[25] M. Correll, D. Albers, S. Franconeri, and M. Gleicher. Comparing averages in time series data. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1095–1104, 2012. 2

[26] Y. Cui, W. G. Lily, Y. Ding, F. Yang, L. Harrison, and M. Kay. Adaptive assessment of visualization literacy. *IEEE Transactions on Visualization and Computer Graphics*, 30(1):628–637, 2023. 2, 3

[27] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, 2019. 7

[28] Y. Ding, J. Wilburn, H. Shrestha, A. Ndlovu, K. Gadhave, C. Nobre, A. Lex, and L. Harrison. revisit: Supporting scalable evaluation of interactive visualizations. In *2023 IEEE Visualization and Visual Analytics (VIS)*, pp. 31–35, 2023. 3

[29] D. Donohoe and E. Costello. Data visualisation literacy in higher education: An exploratory study of understanding of a learning dashboard tool. *International Journal of Emerging Technologies in Learning (iJET)*, 15(17):115–126, 2020. 2

[30] S. E. Dreyfus. The five-stage model of adult skill acquisition. *Bulletin of Science, Technology & Society*, 24(3):177–181, 2004. 2, 7

[31] M. Forehand. Bloom's taxonomy. *Emerging Perspectives on Learning, Teaching, and Technology*, 41(4):47–56, 2010. 2

[32] C. Fosco, V. Casser, A. K. Bedi, P. O'Donovan, A. Hertzmann, and Z. Bylinskii. Predicting visual importance across graphic design types. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*, pp. 249–260, 2020. 2

[33] S. L. Franconeri, L. M. Padilla, P. Shah, J. M. Zacks, and J. Hullman. The science of visual data communication: What works. *Psychological Science in the public interest*, 22(3):110–161, 2021. 2

[34] R. Garcia-Retamero, E. T. Cokely, S. Ghazal, and A. Joeris. Measuring graph literacy without a test: A brief subjective assessment. *Medical Decision Making*, 36(7):854–867, 2016. 1, 2, 3, 4

[35] L. W. Ge, Y. Cui, and M. Kay. CALVI: Critical thinking assessment for literacy in visualizations. In *Proceedings of the 2023 CHI conference on human factors in computing systems*, pp. 1–18, 2023. 1, 2, 3, 4, 5, 8

[36] S. R. Gomez, R. Jianu, R. Cabeen, H. Guo, and D. H. Laidlaw. Fauxvea: Crowdsourcing gaze location estimates for visualization analysis tasks. *IEEE Transactions on Visualization and Computer Graphics*, 23(2):1042–1055, 2016. 2

[37] J. Hong, C. Seto, A. Fan, and R. Maciejewski. Do LLMs have visualization literacy? an evaluation on modified visualizations to test generalization in data interpretation. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–13, 2025. 2

[38] S. Huron, Y. Jansen, and S. Carpendale. Constructing visual representations: Investigating the use of tangible tokens. *IEEE Transactions on Visualization and Computer Graphics*, 20:1, 08 2014. 2

[39] L. Itti, C. Koch, and E. Niebur. A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(11):1254–1259, 1998. 2

[40] N. W. Kim, Z. Bylinskii, M. A. Borkin, K. Z. Gajos, A. Oliva, F. Durand, and H. Pfister. BubbleView: an interface for crowdsourcing image importance maps and tracking visual attention. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 24(5):36, 2017. 1, 2, 3, 5, 9

[41] C. Knittel, J. Awuah, S. Franconeri, and C. X. Bearfield. Gridlines mitigate sine illusion in line charts. In *2024 IEEE Visualization and Visual Analytics (VIS)*, pp. 246–250. IEEE, 2024. 2

[42] J. Kruger and D. Dunning. Unskilled and unaware of it: how difficulties in recognizing one's own incompetence lead to inflated self-assessments.

Journal of Personality and Social Psychology, 77(6):1121, 1999. 4

[43] S. Kullback and R. A. Leibler. *On information and sufficiency*, vol. 22. Institute of Mathematical Statistics, 1951. 6

[44] K. Kurzhals, B. Fisher, M. Burch, and D. Weiskopf. Evaluating visual analytics with eye tracking. In *Proceedings of the Fifth Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization*, pp. 61–69, 2014. 2

[45] B. C. Kwon and B. Lee. A comparative evaluation on online learning approaches using parallel coordinate visualization. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pp. 993–997, 2016. 2

[46] 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. 2, 3

[47] S. Lee, B. C. Kwon, J. Yang, B. C. Lee, and S.-H. Kim. The correlation between users' cognitive characteristics and visualization literacy. *Applied Sciences*, 9(3):488, 2019. 2

[48] E. Lee-Robbins and E. Adar. Affective learning objectives for communicative visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):1–11, 2022. 1

[49] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *IEEE/CVF international conference on computer vision*, 2021. 7

[50] J. Lou, H. Lin, D. Marshall, D. Saupe, and H. Liu. TranSalNet: Towards perceptually relevant visual saliency prediction. *Neurocomputing*, 2022. 2, 7

[51] A. V. Maltese, J. A. Harsh, and D. Svetina. Data visualization literacy: Investigating data interpretation along the novice–expert continuum. *Journal of College Science Teaching*, 45(1):84–90, 2015. 2

[52] 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. 2

[53] L. E. Matzen, M. J. Haass, K. M. Divis, Z. Wang, and A. T. Wilson. Data visualization saliency model: A tool for evaluating abstract data visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):563–573, 2017. 1, 2

[54] P. K. Mudrakarta, A. Taly, M. Sundararajan, and K. Dhamdhere. Did the model understand the question? *arXiv preprint arXiv:1805.05492*, 2018. 6, 9

[55] P. Müller, S. Staal, M. Bâce, and A. Bulling. Designing for noticeability: Understanding the impact of visual importance on desktop notifications. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pp. 1–13, 2022. 1

[56] National Research Council. *How People Learn: Brain, Mind, Experience, and School: Expanded Edition*. National Academies Press, 2000. 1, 3

[57] A. Newman, B. McNamara, C. Fosco, Y. B. Zhang, P. Sukhum, M. Tancik, N. W. Kim, and Z. Bylinskii. Turkeyes: A web-based toolbox for crowdsourcing attention data. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 2020. 3

[58] C. Nobre, K. Zhu, E. Mörth, H. Pfister, and J. Beyer. Reading between the pixels: Investigating the barriers to visualization literacy. In *Proceedings of 2024 CHI Conference on Human Factors in Computing Systems*, 2024. 1, 2

[59] P. O'Donovan, A. Agarwala, and A. Hertzmann. Learning Layouts for Single-PageGraphic Designs. *IEEE Transactions on Visualization and Computer Graphics*, 20(8):1200–1213, Aug. 2014. 2, 3

[60] A. J. Onwuegbuzie and M. A. Seaman. The effect of time constraints and statistics test anxiety on test performance in a statistics course. *The Journal of Experimental Education*, 63(2):115–124, 1995. 9

[61] S. Pandey. Benchmarking visual language models on standardized visualization literacy tests. In *Proceedings of the Eurographics*, 2025. 2

[62] S. Pandey and A. Ottley. Mini-VLAT: A short and effective measure of visualization literacy. In *Computer Graphics Forum*, vol. 42, pp. 1–11. Wiley Online Library, 2023. 1, 2, 3, 4

[63] A. Papoutsaki, P. Sangkloy, J. Laskey, N. Daskalova, J. Huang, and J. Hays. WebGazer: Scalable webcam eye tracking using user interactions. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 3839–3845. AAAI, 2016. 2, 3, 9

[64] E. M. Peck, S. E. Ayuso, and O. El-Etr. Data is personal: Attitudes and perceptions of data visualization in rural pennsylvania. In *Proceedings of the 2019 CHI conference on human factors in computing systems*, 2019. 1

[65] A. M. Persky and J. D. Robinson. Moving from novice to expertise and its implications for instruction. *American Journal of Pharmaceutical*

Education, 81(9):6065, 2017. 2, 7

[66] P. Polatsek, M. Waldner, I. Viola, P. Kapec, and W. Benesova. Exploring visual attention and saliency modeling for task-based visual analysis. *Computers & Graphics*, 72:26–38, 2018. 2

[67] G. J. Quadri, A. Z. Wang, Z. Wang, J. Adorno, P. Rosen, and D. A. Szafir. Do you see what i see? a qualitative study eliciting high-level visualization comprehension. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, pp. 1–26, 2024. 1, 2

[68] T. Reding and C. Moore. How expert and inexpert instructors talk about teaching. *Education Sciences*, 13(6):591, 2023. 2, 6

[69] W. Revelle and M. W. Revelle. Package 'psych'. *The Comprehensive R Archive Network*, 337(338):161–165, 2015. 4

[70] A. M. B. Rodrigues, G. D. J. Barbosa, H. C. V. Lopes, and S. D. J. Barbosa. What questions reveal about novices' attempts to make sense of data visualizations: Patterns and misconceptions. *Computers & Graphics*, 94:32–42, 2021. 2

[71] P. Ruchikachorn and K. Mueller. Learning visualizations by analogy: Promoting visual literacy through visualization morphing. *IEEE Transactions on Visualization and Computer Graphics*, 21(9):1028–1044, 2015. 2

[72] D. Rudoy, D. B. Goldman, E. Shechtman, and L. Zelnik-Manor. Crowdsourcing gaze data collection. *CoRR*, abs/1204.3367, 2012. 2, 3

[73] K. J. Schönborn and T. R. Anderson. The importance of visual literacy in the education of biochemists. *Biochemistry and Molecular Biology Education*, 34(2):94–102, 2006. 2

[74] K. Shanmuga Vadivel, T. Ngo, M. Eckstein, and B. Manjunath. Eye tracking assisted extraction of attentionally important objects from videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3241–3250, 2015. 2

[75] H. Shin, Q. Zou, and W. J. Ma. The effects of delay duration on visual working memory for orientation. *Journal of vision*, 17(14):10–10, 2017. 9

[76] S. Shin, S. Chung, S. Hong, and N. Elmqvist. A scanner deeply: Predicting gaze heatmaps on visualizations using crowdsourced eye movement data. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):396–406, 2022. 1, 2

[77] C. Spearman. The proof and measurement of association between two things. *The American Journal of Psychology*, 15(1):72–101, 1904. 6

[78] H. Suresh, S. R. Gomez, K. K. Nam, and A. Satyanarayan. Beyond expertise and roles: A framework to characterize the stakeholders of interpretable machine learning and their needs. In *Proceedings of the 2021 CHI conference on human factors in computing systems*, pp. 1–16, 2021. 1

[79] D. Toker, C. Conati, B. Steichen, and G. Carenini. Individual user characteristics and information visualization: connecting the dots through eye tracking. In *proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 295–304, 2013. 1, 3

[80] Y. Tong, H. Konik, F. Cheikh, and A. Tremeau. Full reference image quality assessment based on saliency map analysis. *Journal of Imaging Science and Technology*, 54(3):30503–1, 2010. 1, 6

[81] A. M. Treisman and G. Gelade. A feature-integration theory of attention. *Cognitive psychology*, 12(1):97–136, 1980. 2

[82] A. van den Brandt, S. L'Yi, H. N. Nguyen, A. Vilanova, and N. Gehlenborg. Understanding visualization authoring techniques for genomics data in the context of personas and tasks. *IEEE Transactions on Visualization and Computer Graphics*, 2024. 3

[83] P. Viola and W. M. Wells III. Alignment by maximization of mutual information. *International Journal of Computer Vision*, 24(2), 1997. 6

[84] Y. Wang, M. Bâce, and A. Bulling. Scanpath prediction on information visualisations. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 30(7):3902–3914, 2024. 2

[85] Y. Wang, M. Koch, M. Bâce, D. Weiskopf, and A. Bulling. Impact of gaze uncertainty on aois in information visualisations. In *2022 symposium on eye tracking research and applications*, pp. 1–6, 2022. 1

[86] Y. Wang, W. Wang, A. Abdelhafez, M. Elfares, Z. Hu, M. Bâce, and A. Bulling. SalChartQA: Question-driven saliency on information visualisations. In *Proc. ACM SIGCHI Conference on Human Factors in Computing Systems (CHI)*, pp. 1–14, 2024. 2, 3, 5, 6, 7

[87] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004. 6

[88] Z. Wang, L. Sundin, D. Murray-Rust, and B. Bach. Cheat sheets for data visualization techniques. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pp. 1–13, 2020. 2