Reading Between the Pixels: Investigating the Barriers to Visualization Literacy

Carolina Nobre
cnobre@cs.toronto.edu
University of Toronto
Toronto, Canada

Kehang Zhu
Harvard University
Boston, USA
kehang_zhu@g.harvard.edu

Eric Mörth
Harvard University
Boston, USA
ericmoerth@g.harvard.edu

Hanspeter Pfister
Harvard University
Boston, USA
hpfister@harvard.edu

Johanna Beyer
Harvard University
Boston, USA
jbeyer@harvard.edu

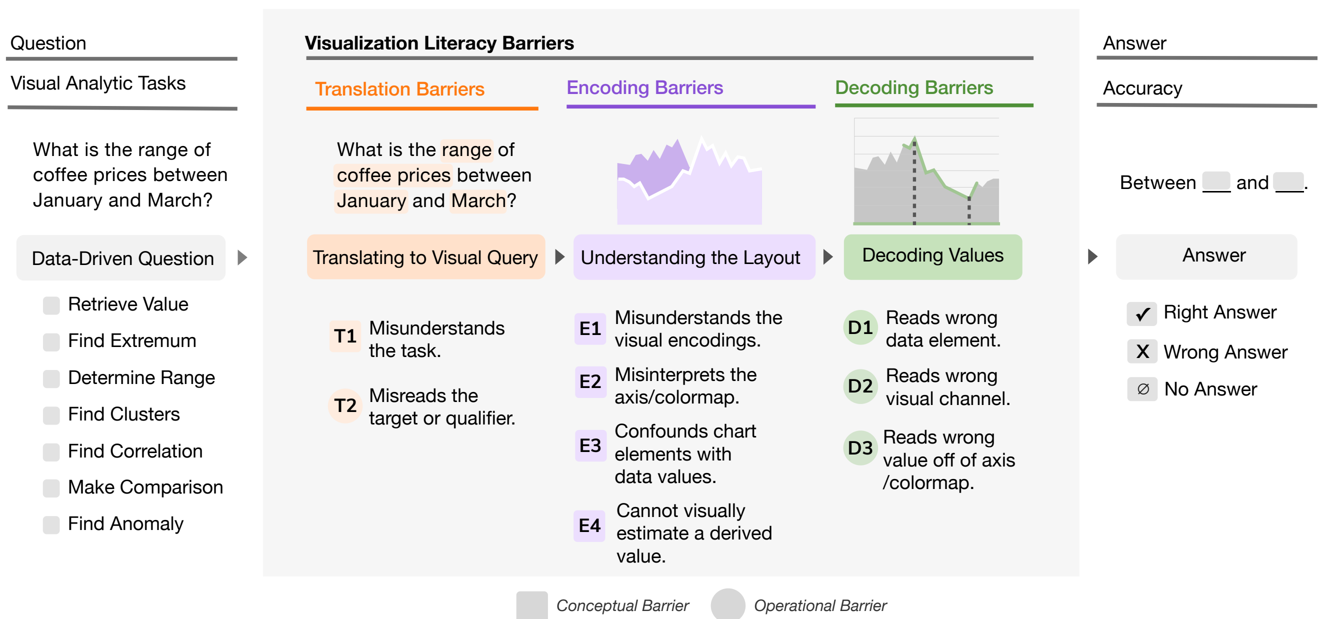

Figure 1: Overview of Barriers to Visualization Literacy. An overview of barriers at the various stages of reading a data visualization. The leftmost panel delineates various visualization tasks. The central panel outlines the three primary barriers faced during interpretation: translation, encoding, and decoding. Barriers can either be *conceptual*, i.e., resulting from a flaw in understanding, or *operational*. i.e., resulting from a mistake in reading a value in the visualization or term in the question. *Translation barriers* are those that occur when extracting the task and target of a question. *Encoding barriers* are those that result from a misunderstanding of the visual encodings in the chart. *Decoding barriers* capture the inability to properly extract a data value from the chart. The rightmost panel showcases the potential outcomes: a correct answer, an incorrect answer, or no answer.

ABSTRACT

In our current visual-centric digital age, the capability to interpret, understand, and produce visual representations of data —termed visualization literacy— is paramount. However, not everyone is adept at navigating this visual terrain. This paper explores the barriers that individuals who misread a visualization encounter, aiming to understand their specific mental gaps.

Utilizing a mixed-method approach, we administered the Visualization Literacy Assessment Test (VLAT) to a group of 120 participants drawn from diverse demographic backgrounds, which

CHI '24, May 11–16, 2024, Honolulu, HI, USA

ACM ISBN 979-8-4007-0330-0/24/05
https://doi.org/10.1145/3613904.3642760

provided us with 1774 task completions. We augmented the standard VLAT test to capture quantitative and qualitative data on participants' errors. We collected participant sketches and open-ended text about their analysis approach, providing insight into users' mental models and rationale.

Our findings reveal that individuals who incorrectly answer visualization literacy questions often misread visual channels, confound chart labels with data values, or struggle to translate data-driven questions into visual queries. Recognizing and bridging visualization literacy gaps not only ensures inclusivity but also enhances the overall effectiveness of visual communication in our society.

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in visualization**; *Visualization theory, concepts and paradigms.*

KEYWORDS

Visualization, visualization literacy, conceptual barriers

ACM Reference Format:

Carolina Nobre, Kehang Zhu, Eric Mörth, Hanspeter Pfister, and Johanna Beyer. 2024. Reading Between the Pixels: Investigating the Barriers to Visualization Literacy. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24), May 11–16, 2024, Honolulu, HI, USA.* ACM, New York, NY, USA, 17 pages. https://doi.org/10.1145/3613904.3642760

1 INTRODUCTION

In the age of Instagram, infographics, and interactive dashboards, visual data permeates our daily lives. Every swipe on our smartphones, every website we visit, all resonate with visual cues designed to capture attention and convey messages. This omnipresence of visual stimuli in our digital age accentuates the importance of visual literacy—the ability to decode, interpret, and make sense of visual information [4]. *Visualization literacy*, in particular, describes a person's ability to understand data presented in graphical elements, such as charts, graphs, and maps. Just as traditional literacy paved the way for success in the text-dominant eras of the past, visualization literacy is the key to unlocking the potential of our increasingly visual-centric world.

However, not all individuals sail smoothly on this sea of visuals. Just as with textual literacy, gaps exist, and these disparities in visualization literacy can lead to significant misunderstandings and misinterpretations of data. Understanding the nature and depth of these gaps is not just a pedagogical concern but also a social one. It touches upon the very essence of how information is received, processed, and acted upon by diverse sections of society.

The task of measuring and assessing visual literacy, especially in pedagogical contexts, has garnered heightened interest over the years [4, 22]. More recently, tests for evaluating an individual's visualization literacy have been developed, such as the VLAT [26] or Mini-VLAT [28]. These tests use a multitude of different visualization types and questions to arrive at a numerical score of visualization literacy for the test taker.

While extensive research has been dedicated to the development and assessment of visualization literacy tests [15, 19, 26, 28], these tests do not inherently offer insights into the root causes of visualization illiteracy. Therefore, we posit that we have to go beyond mere testing to identify strategies for enhancing visualization literacy. It is imperative to first understand the underlying misconceptions users harbor when interpreting visualizations.

In this paper, we investigate some of the barriers to visualization literacy and look at the challenges faced by those who grapple with visual representations. We leverage existing tools that measure viewers' ability to perform visual tasks with basic charts (the VLAT) and focus on understanding the mental gaps and conceptual misunderstandings that can occur. We do this through a large crowdsourced study that captures both quantitative and rich qualitative data on participants who take the VLAT. With this, we aim to shed light on the specific hurdles, nuances, and potential pitfalls that individuals encounter when they misread a visualization. From our research, we introduce a classification of challenges that viewers can encounter when making sense of a visualization. We classify obstacles into a) difficulties in translating a question to a visual query (translation barriers), b) difficulties in understanding the layout and visual encoding of a visualization (encoding barriers), and c) mistakes in reading out and decoding values of a visualization (decoding barriers). We elucidate each barrier with examples from our study.

In this work, we make two main contributions:

- A mixed-methods empirical study that investigates the rationale behind mistakes in the Visualization Literacy Assessment Test.
- A classification of the different barriers in interpreting data visualizations.

2 RELATED WORK

Visualization literacy has been at the forefront of recent academic investigations, with scholars striving to understand how individuals interpret and derive meaning from visual data representations. In this section, we explore the related work in the field of visualization literacy, including definitions, assessment methods, and conceptual barriers.

2.1 Visualization literacy definition and measurement

The concept of visualization literacy has its roots in the early 20th century. The broader term *visual literacy* was popularized in the 1960s and 1970s as the ability to understand, interpret, and produce visual messages [11, 16, 31, 33, 35]. More specific to the domain of data and information visualization, the term *visualization literacy* evolved several years later, as the ability to interpret (read) and create (write) data visualizations effectively. More specifically, Börner et al. defined visualization literacy as: "The ability to make meaning from and interpret patterns, trends, and correlations in visual representations of data [8].

Offering a parallel perspective, Lee et al. describe it as "the aptitude and capability to read, interpret, and glean information from data visualizations." [25]. In this study, we adopt Lee et al.'s definition, as it aligns closely with the skills evaluated by the Visualization Literacy Assessment Test (VLAT) we utilized.

In a distinct approach to gauge visualization literacy, Boy et al. [10] utilized the response theory to develop measurement items. Their study identified six key tasks for visualization comprehension:

determining minimum and maximum values, discerning variation, pinpointing intersections, computing averages, and making comparisons. Their evaluation tools incorporated line charts, bar charts, and scatterplots.

Building on the foundational work of Boy et al., Lee et al., and Borner et al. in defining visualization literacy, research progressively expanded to develop effective methods of assessing it. To this end, Lee et al. [26] introduced the now widely recognized and established Visualization Literacy Assessment Test (VLAT), comprising 12 data visualizations and 53 multiple-choice items. The visualizations include the following: line chart, bar chart, stacked bar chart, 100% stacked bar chart, pie chart, histogram, scatterplot, bubble chart, area chart, stacked area chart, choropleth map, and treemap. The test featured high validity and reliability, but even with a limit of 25 seconds per question, the VLAT is approximately 22 minutes long. This motivated Pandey et al.[28] to develop a shortened version of the VLAT with only 12 items instead of 53, which is still reliable and strongly correlating with the original VLAT. Several other approaches to test visualization literacy have been proposed recently[15], often requiring less time than the VLAT test. Firat et al. [19] present a survey on visualization literacy and its evaluation. In this study, we employ the full VLAT survey by Lee et al. to prompt participants to perform visual analytic tasks with basic charts. However, our main goal is not to assess or improve VLAT scores but to understand the mental models that lead to incorrect answers. Therefore, we also collect open-ended responses that provide us with qualitative data to better understand participant rationale.

2.2 Understanding Barriers to Visualization Literacy

While much of the literature concentrates on measuring and enhancing visualization literacy, fewer studies have dissected the specific conceptual obstacles and mental models contributing to limited visualization literacy. In their state of the art report on interactive visualization literacy, Firat et al. [18] note that the most common future research goal identified in their corpus of papers was 'an improved understanding of barriers to visualization literacy'

In one such effort, Grammel et al. [21] studied the behaviors of novice users as they created visualizations from a given dataset. They identified several challenges, which they termed "barriers" faced during this process. These barriers included high visual complexity, unfamiliar visualization types, difficulties understanding the semantics of measurements, and readability problems. Although these barriers offer valuable insights, they remain broad and do not delve into the specific aspects of the visualization that pose challenges for the viewer.

In a similar vein, Kwon et al. [13] delved into the difficulties, or "roadblocks" as they called them, that novice users encounter in investigative analysis using the mature visualization tool, Jigsaw. The roadblocks encountered mostly related to the interactive nature of the target tool, such as "Failure to execute appropriate interactions" or very high-level abstractions, such as "Failure to interpret visualizations".

In contrast, Börner et al. [8] investigated cognitive barriers to visualization literacy with static visualizations of varying levels of complexity. This investigation involved the presentation of a range of visualization types within the premises of three science museums in the United States, coupled with inquiries directed at visitors to gauge their knowledge of the displayed visualizations. Through this study, the authors sought to elucidate the cognitive boundaries associated with individuals' capacity to recognize and comprehend various forms of data visualizations. The findings showed that individuals possess the capability to comprehend and label visualizations only if they have prior familiarity with them. In the case of more intricate visualizations, such as graph layouts, visitors encountered difficulties in deciphering the underlying data representations conveyed by the presented visuals.

With a focus on a specific chart type, Peebles et al. [29] explored how those unfamiliar with parallel-coordinates plots interpret such visualizations. They identified two primary impediments to accurate interpretation: (1) a lack of understanding of coordinate systems and (2) visual complications from crossing lines and clutter. In contrast to previous studies, the authors highlighted specific design elements that caused confusion. Our work further delves into these barriers to visualization literacy, examining potential hurdles across a broader array of visualization types.

Thompson et al. [32] conducted a comprehensive investigation by engaging with experts in the fields of visualization literacy and information literacy. Their study aimed to discern emerging patterns, obstacles, and prospects that are shaping the landscape of visualization literacy in the twenty-first century. Their particular emphasis was on the aptitude to interpret and analyze visual content, revealing a notable deficiency in addressing what is termed as "visualization literacy" within the United States' educational framework [3]. Furthermore, Thompson et al. observed discernible socioeconomic disparities among students concerning their proficiency in visualization literacy. Meanwhile, in research focused on visualization literacy barriers in graph exploration, Alkadi et al. [2] investigated challenges faced by users in visualizing networks. A key insight from their work was that users frequently faced difficulties when their analyses clashed with pre-existing mental models or expectations of how their networks should appear.

Satkowski et al. [30] showed that individuals with lower visualization literacy might benefit more from adapted visualizations (i.e., visualizations that have been modified to highlight important elements) than individuals with high visualization literacy. Similarly, Birchfield et al. [7] found that feedback mechanisms improve graphical perception when reading charts. Finally, Yang et al. [37] explored how individuals explain visualizations. However, none of these studies specifically investigated the different barriers to visualization literacy.

Our work is closely connected to the research conducted by Lee et al. [27]. In their investigation, the authors delved into the correlation between visualization literacy and individual factors such as numeracy, need for cognition, and cognitive style classified as visualizer-verbalizer. Their findings revealed a statistically significant positive correlation between numeracy and visualization literacy, as well as a positive correlation with the need for cognition. Notably, no discernible correlation emerged between the visualizer-verbalizer cognitive style and visualization literacy. These identified correlations offer empirical insights into the nuanced barriers that individuals encounter in attaining visualization literacy. Moreover, they provide a foundational basis for informing subsequent studies and guiding the development of targeted interventions to address specific barriers on an individualized level. Our work focuses less on the individual factors that contribute to visualization literacy and more on identifying general barriers to visualization literacy.

Boy et al. [9] delineated four constituent sub-costs within Wijk's framework of perception and exploration costs (Ce) [36]: 1) literacy cost, 2) context-interpretation cost, 3) perceived interactivity cost, and 4) initial incentive for exploration cost. Cost in the presented framework can also be seen as a barrier. If a certain individual threshold is reached, the cost is too high in comparison to the benefits, manifesting a barrier for the individual visualization reader. The first cost characterizes the challenge encountered by information seekers unaccustomed to extracting meaning from visualizations. The second cost pertains to the cognitive effort required to contextualize information within a visualization, encompassing the interpretation of titles, labels, annotations, and other elements. The perceived interactivity cost (3) elucidates a heightened cost/benefit ratio associated with the identification of interactive features within a graphic. Finally, a deficiency in an information seeker's background knowledge concerning the visualized dataset or the indicators employed for information representation results in a diminished motivation to engage in exploratory data analysis. The authors provide a nuanced understanding of the challenges inherent in the interaction with visualized data. The presented costs offer valuable insights to inform strategies aimed at better understanding potential barriers in visualization literacy.

Existing literature on barriers to visualization literacy reveals a need for a more granular comprehension of conceptual barriers and how they manifest across different chart types. Our work addresses this by collecting in-depth qualitative data on the thought processes and reasoning behind them. We examined the obstacles that arise at every stage of understanding and extracting information from visual representations. Additionally, we delve into specific misconceptions and focus on the difficulties people face when trying to comprehend various types of visualizations. Finally, our study encompasses twelve different chart types, providing evidence for the misconceptions that occur with each type. From this research, we have developed a classification that outlines the barriers encountered at each stage of the process of extracting information from data visualizations.

3 CROWDSOURCED STUDY

3.1 Study Design

We designed our study to pinpoint the challenges participants faced while completing visualization literacy tasks. We employed the 53 questions from the VLAT survey [26], which encompasses the 12 visualization types outlined in Figure 3. The chart designs are from the more recent version of the VLAT by Pandey et al.[28]. To better understand the conceptual flaws that can lead viewers to incorrectly extract information from a visualization, we made two main changes to the way in which VLAT questions were presented to the viewer: (1) removing multiple choice options, and (2) adding in qualitative data collection.

To ensure a wide range of responses, we removed the multiple-choice options from the VLAT questions, except in cases with a limited set of possible answers, like specific US states or a finite set of data categories. When questions sought numerical responses, we had participants input an open-ended number. This approach not only minimizes guesswork but also sheds light on unique, often unconventional, cognitive frameworks. Additionally, we provided an option for participants to indicate uncertainty with an "I'm not sure" response.

In our second modification to the VLAT questions, we collected two types of qualitative data from participants: (1) free-form annotations on the visualizations, and (2) open-ended text that described their reasoning during the task. We used the Qualtrics platform to conduct the study, which allowed us to implement display logic to request sketches and explanations only when participants either gave an incorrect answer or chose "I'm not sure", thus minimizing the overall cognitive load of taking the survey.

For qualitative feedback, we displayed the original VLAT question next to the participant's answer and asked them to annotate directly on the visualization with their mouse. We further prompted them with: "Can you explain your thought process? What elements guided your decision? Did you find any part confusing?"

3.2 Pilots and Experiment Planning

We conducted three pilots, with 30, 60, and 60 participants, respectively, to evaluate tasks, visualization design, and our procedure. Initial pilots revealed ambiguity in the design of some of the VLAT charts from Pandey et al.[28]. For example, the x-axis labels for time-based data did not clarify whether a year label was the start or end of the year, creating confusion among participants. Similarly, the treemap design, which did not include a box around all categories, left several participants unclear as to whether they all added up to 100%. Since the goal of this study was to pinpoint conceptual gaps in viewers' understanding of the visualization, we wanted to minimize the effect of mistakes that resulted from the poor or ambiguous design of the visualizations themselves. To this end, we updated the design of the charts to make the labeling of the axis and the treemap design as clear as possible. These modifications also align with the design of these charts in the original VLAT survey[26].

Another chart design insight that emerged from the pilots related to the labeling of bars in the histogram. The original design of the VLAT histogram, as well as most common histograms, labels the ranges for each bin at the tick marks on the right and left side of the bar (Figure 2, left panel). However, several participants in the pilot interpreted the bar as representing the count for a single x value, instead of the range between the two boundary values. To further investigate this conceptual ambiguity, we created and used two designs for the histogram in the final study, one where the x ticks/labels on the edges of the bar represent the boundary values (Figure 2 left panel), and the other where the label is directly under the bar and shows the exact range for that bar (e.g 10-20) (Figure 2 right panel). Participants who saw a histogram were randomly assigned one of the two histogram designs, and the analysis of the results was done for each histogram separately (Table 2). Results show that participants performed better with the updated design, which reduced the ambiguity of what data was contained in each bar.

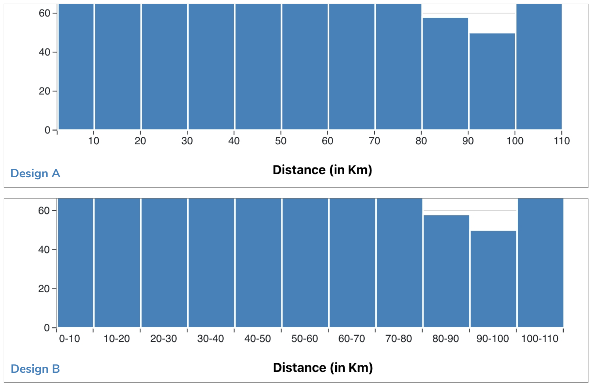

Figure 2: Histogram x-axis labeling strategies. The top panel represents the more standard x labeling for histograms, where bars sit in between the ticks that denote the range of values contained in that bar. The bottom panel shows an alternate design, which centrally positions a single label beneath each bar to denote its range. Study findings revealed superior participant performance with the bottom design compared to the top.

Data from the pilots also revealed that participants frequently mistook the first and last data points – as opposed to the max and min – as the data range in the area chart. To determine if the data trend in the chart (with the maximum value at the start and minimum one data point shy of the end) played a role in this confusion, we devised an alternative area cart that presented a markedly different data pattern (see inset in the Area chart in Figure 3). Participants performed better with the alternative area chart in all but the retrieve value task, where the accuracy was the same between both datasets.

Lastly, since we were focused on understanding the rationale behind incorrect task completions, the final study did not include the four questions that all 150 pilot participants answered correctly. Specifically, these tasks were bar chart questions 1, 2, and 3 and Choropleth question 3.

3.3 Participants and Procedure

We recruited 120 participants for the full study on Prolific, a crowdsourcing platform with a research focus. The participants included 56 men, 64 women, and 2 non-binary/other. Participant ages were between 18 and 76. We also collected participants' self-reported familiarity with data visualizations, which included no familiarity (5), beginner (35), moderately knowledgeable (56), relatively skilled (25), and expert (1). A complete breakdown of participant demographic information can be found in our supplementary material. Our study used a mixed-subjects design in which each participant was randomly assigned to four of the twelve VLAT charts. The twelve charts used can be seen in Figure 3. We restricted the number of charts to 4 to avoid survey fatigue, particularly with the added cognitive load of providing sketches and explanations. With this design, we ended up with 40 participant entries for each of the 12 charts. Since each chart has between 3 and 6 questions in the VLAT, we collected a total of 1774 task completions. Of these, 490 were incorrect (27%) and were accompanied by sketches and explanations. We deployed the study on the Qualtrics survey platform, which ensures random distribution amongst conditions.

The study consisted of three sections: *Tutorial*, *VLAT questions*, and *Demographics*.

The tutorial section at the start of the study introduced participants to the format of VLAT questions, as well as how to sketch on the visualizations and answer the open-ended rationale questions. After the tutorial, participants answered the VLAT questions for their randomly assigned chart types. All questions included the option for participants to respond 'I'm not sure' or 'Impossible to answer with this visualization'. When participants answered the task incorrectly, we elicited the follow-up sketch and open text question. The last section collected demographic information, as well as their self-assessed experience with visualization and whether they have any sort of color blindness. The complete survey is available in the supplementary material.

Based on the completion times of the pilot experiments, each participant was paid $4 USD, for an estimated duration of 20 minutes, resulting in an hourly rate of about $12 USD. The median time of completion after the survey was completed was 25 minutes. All participants viewed and agreed to an IRB-approved consent form. To be eligible for the study, participants had to use a laptop or desktop device with a resolution of at least 1400x850 pixels available screen space in the browser.

4 A CLASSIFICATION OF BARRIERS TO VISUALIZATION LITERACY

The results from the study provided two lenses for understanding barriers to visualization literacy for the different chart types. The *qualitative* data includes the participant sketches as well as written explanations in text. The analysis of this data led to the classification of barriers to visualization literacy as shown in Figure 1 and is described in detail in this section. The *quantitative* data, which includes accuracy and time, revealed the charts and tasks that participants struggled with the most. Analyzing the performance and time metrics in light of the barriers identified from the qualitative data provided us with rich insight into the main barriers associated with different tasks and charts. We present and discuss these results in Section 5.

In this section, we describe how we analyzed the study data to produce a categorization of barriers in reading a visualization. We delineate the sequential steps that span from reading the question to formulating an answer, highlighting key barriers at each juncture. Drawing from our survey responses, we present examples of these common pitfalls.

4.1 Methodology

To shed light on participants' underlying cognitive barriers and mental models, we performed a qualitative analysis of all the sketches and open-form text in incorrect task completions. For this, we employed the open coding method, which involves systematically identifying and tagging patterns, themes, or concepts within the collected data [14]. This process was done by both the first and

Figure 3: Chart Designs. The 12 charts used in the crowdsourced survey: area chart, bar chart, line chart, pie chart, choropleth chart, stacked bar chart 100%, scatter plot, treemap, histogram chart, bubble chart, stacked area chart, stacked bar chart. Two charts had dual designs, which emerged from findings in our pilot. The histogram variants had different x-axis labeling (see inset). The area chart variants each had a different data pattern (see inset). Color palettes were chosen to ensure they were colorblind-safe.

second authors of the paper independently, with any discrepancies between them resolved on a case-by-case basis. This method ensured a comprehensive and unbiased understanding of how participants navigated and decoded the visualizations presented to them. The final set of codes can be found in the supplementary material.

The output of the coding process by the authors was then organized into higher-level categories of different types of mistakes. For example, all the codes that captured a mistake in the process of translating the question to a visualization task were grouped into a 'translation' category. This systematic grouping of codes into higher-level groups led to the classification of barriers presented in Figure 1, which form the main contribution of this work.

4.2 Classification of Barriers

The process of answering a question with a data visualization can be broken down into 5 components, as shown in Figure 1:

(1) Posing the question
(2) Translating the question into a visual query,
(3) Grasping the visualization's layout and visual encodings,
(4) Decoding relevant values, and
(5) Synthesizing an answer.

Barriers that occur during stages 2, 3, and 4 are classified as *translation*, *encoding*, and *decoding* barriers, respectively. Our study design is not well suited to characterize barriers specific to stages 1 and 5 since we employ pre-defined tasks (VLAT) which have specific correct answers. Barriers in these stages would be best explored in a study with open-ended exploratory tasks. For the three types of barriers we characterize in this work, we further distinguish between *conceptual* and *operational* barriers. *Conceptual barriers* are those that result from a gap in the viewer's *understanding* of the task or the visual encodings used. For example, if the viewer does not know the meaning of a term in the task, or they do not understand how a visual variable maps to the underlying data. *Operational barriers*, on the other hand, are those that occur when a user misreads the question or elements in the visualization, such as data values or color legends. Operational barriers can result from a previous conceptual barrier, but can also happen in isolation as a result of inattention. We elaborate further on each type of barrier in the subsequent sections.

4.3 Translation Barriers: From Text to Visual Query

During the translation phase, users interpret the question and discern what information they need to derive from the visualization. Mistakes in this phase can be categorized into two main types:

- T1: Misunderstands the visual task.
- T2: Misreads the target or qualifier of the task.

Errors during translation often cascade to subsequent stages, possibly leading to mistakes like decoding the wrong data elements (D1) or accessing the incorrect visual channel (D2). An in-depth discussion on decoding errors can be found in Section 4.5.

4.3.1 *T1: Misunderstands the task.*

A *conceptual* barrier emerges when the user fails to understand the visualization task requested. Such misunderstandings frequently arise when a question uses terminology unfamiliar to the viewer, like 'range', 'ratio', or 'outlier'.

In our research, many participants misunderstood the task in question. For instance, when asked to determine the "range of coffee prices observed between January 2018 and 2019" (ACQ3), many mistook 'range' for 'difference in price between the two dates'. One participant explained, "I found the values for the two months, then ordered them from first to last." (See Figure 4a for a visual example).

4.3.2 *T2: Misreads the target or qualifier.*

In contrast to T1, this is an *operational* barrier, and stems from the viewer directing their attention to the wrong target. This might manifest as identifying the accurate price for an incorrect category. Causes for this misstep can range from hastily reading the question to inadvertently focusing on a more visually prominent, yet incorrect, target. Survey instances of this error included identifying the right month but in the wrong year in a time series, the wrong category in a Stacked Bar (Figure 4b), or area chart, or an incorrect interval on a histogram.

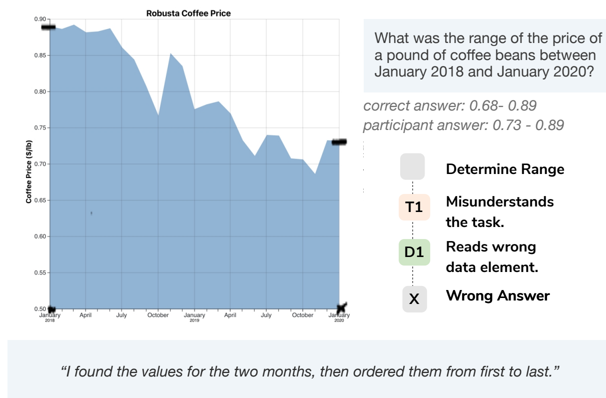

(a) Example of a *T1: Misunderstands the task* mistake. The participant misunderstood what "range' of the data meant and instead read the values only at the start and end dates. This mistake also led to D1 mistake, where the participant read off the wrong data element in the chart.

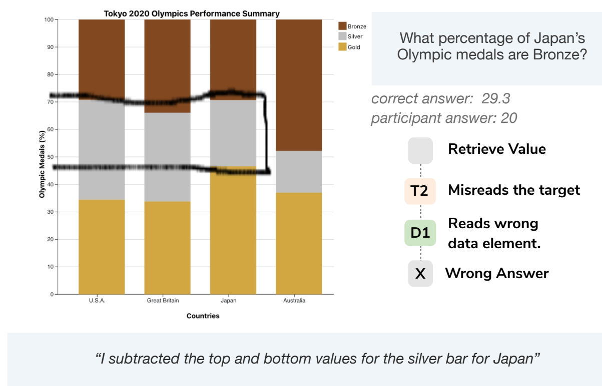

(b) Example of a *T2: Misreads the target/qualifier mistake.* The participant looked up the silver medals for Japan instead of the Bronze as asked in the task. This mistake also cascaded into a D1 decoding mistake, where the participant read off the wrong data element from the chart.

Figure 4: Examples of Translation and Decoding Barriers encountered by participants in the survey. The sketches portray participant annotations, while the pathway on the right of each panel denotes the sequence of mistakes made by the participant.

4.4 Encoding Barriers: Understanding Marks and Channels

Encoding barriers arise when a viewer struggles to comprehend the chart's visual encodings, including the visual channels, structural chart elements, and interpretation of axes and color maps. The barriers in this category are all *conceptual* since they result from a misunderstanding of how the data is mapped to visual variables. Such errors can fall into one of four categories:

- E1: Misunderstands the visual encodings
- E2: Misinterprets the axis or colormap
- E3: Confounds chart elements with data values
- E4: Cannot visually estimate a derived value

Mistakes at this stage can also lead to decoding errors in subsequent steps, such as targeting the wrong data value, or can halt the analysis entirely if the viewer cannot comprehend the visual encodings.

4.4.1 *E1: Misunderstands the visual encodings.*

The most critical of encoding barriers arises when users misinterpret the specific visual cues, such as colors, shapes, or sizes, used to represent data. For instance, a viewer might misinterpret the size of a circle in a bubble chart as an indicator of one data dimension when it actually represents another. A particularly common example of this misconception is assuming the height of each segment in a stacked bar chart reflects the absolute value for its respective category. Such misunderstandings can arise from unfamiliarity with standard encoding practices or when a visualization employs unconventional or complex visual mappings. This type of error most often leads to one of the decoding errors described in Section 4.5.

In our survey, this error was particularly prevalent in stacked area charts, stacked bar charts, and 100% stacked bar charts, where participants interpreted the height of a stacked bar or area as the absolute value for that category. More unique misunderstandings of the layout included a belief that histogram bars were cumulative (Figure 5a) or that stacked bars represented ranges of prices for that item (Figure 5b)

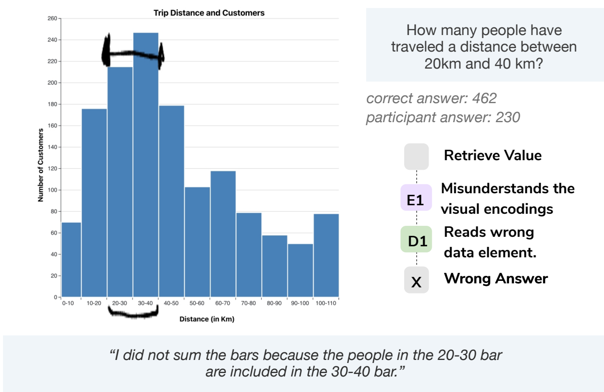

(a) Example of a *E1: Misunderstands the visual encodings* mistake. The participant misunderstood what each bar in the histogram represents. This also led to a D1 decoding mistake, where they read off the values for the wrong bars in the chart.

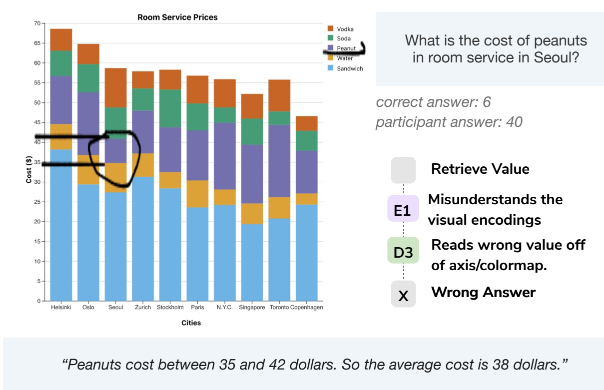

(b) Example of a *E1: Misunderstands the visual encodings* mistake. The participant interpreted stacked bars as representing the min-max range for each category. This conceptual hurdle led to a subsequence D3 mistake, where the participant read off the wrong value from the axis (the middle of the bar).

Figure 5: Examples of Encoding (E1) and Decoding (D1/D3) Barriers encountered by participants in the survey. The sketches portray participant annotations, while the pathway on the right of each panel denotes the sequence of mistakes made by the participant.

4.4.2 *E2: Misinterprets the axis or colormap.*

This obstacle results from a misinterpretation of axis values, particularly when viewers are tasked with deducing values that fall between marked intervals. For example, in a time-series chart, a viewer might infer that the line connecting data points at two tick marks represents values for the individual days that span between the two timestamps.

This particular conceptual error was frequently observed in time series charts in our study (Figure 6a). A similar type of mistake occurred when interpreting the quantized colormap in the Choropleth map. Viewers found the correct color but did not know how to read the value from the quantized colormap legend (Figure 6b).

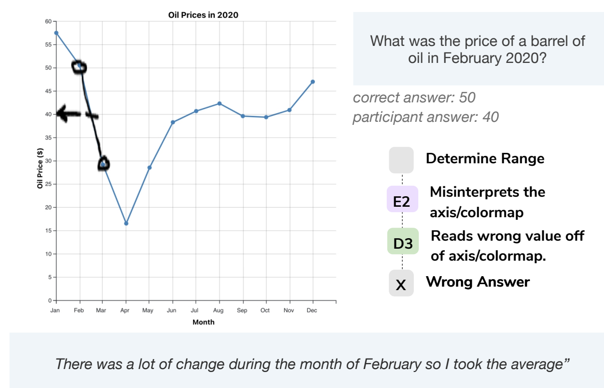

(a) Example of a *E2: Misinterprets the axis or colormap* mistake. The participant interprets the line connecting the Feb and March data points as the days in February. This conceptual flaw also led to the D3 mistake of reading the wrong value off of the Y axis.

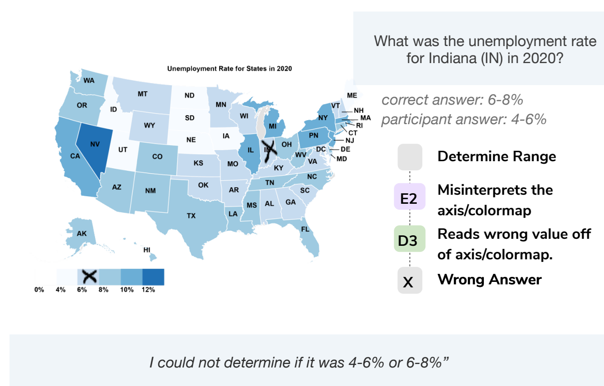

(b) Example of a *E2: Misinterprets the axis or colormap* mistake. The participant correctly identifies the color but does not realize that the color refers to the range 6- 8%. As a result, they also fall victim to the D3 mistake of reading off the wrong value from the colormap.

Figure 6: Examples of Encoding (E2) and Decoding (D3) Barriers encountered by participants in the survey. The sketches portray participant annotations, while the pathway on the right of each panel denotes the sequence of mistakes made by the participant.

4.4.3 *E3: Confounds chart elements with data values.*

This barrier occurs when viewers mistakenly associate structural chart elements, such as axis labels, gridlines, or reference elements, with the data being depicted. For instance, a viewer might misinterpret an axis value or a reference line as part of the dataset, leading to skewed perceptions of values or trends.

During our survey, this mistake appeared when participants were tasked with indicating the value range in the dataset, especially in the Area chart where the data mark extends to the bottom of the axis, but the actual values are denoted at the top of the area. In these situations, many participants bypassed the actual data values and answered based on the minimum and maximum labels on the Y-axis (Figure 7a). Another instance of this mistake was observed when viewers were asked to estimate the average value in the Bubble chart, and they highlighted the center of the chart as being the average value (Figure 7b).

4.4.4 ***E4: Cannot visually estimate a derived value.*** This barrier arises when viewers are faced with the challenge of deducing a value that is not explicitly presented in the visualization, such as an average or a noticeable trend among data points. Viewers may be hesitant about their estimates or uncertain about the methodology to derive the value. In our survey, when faced with questions asking for derived values, participants typically provided one of two incorrect responses. Some expressed their uncertainty by responding with 'I'm not sure', indicating confusion about visually estimating an average. Others responded with 'Impossible to Answer with this Visualization', suggesting the data presented did not allow for an accurate estimation.

4.5 Decoding Barriers: Reading the Visualization

Decoding a visualization involves a viewer's ability to correctly extract the information presented through visual channels. Decoding barriers arise when viewers encounter difficulties in reading off information from the visualization, and are all examples of *operational* barriers. Common mistakes include misreading data values, referencing the wrong axis or dimension, or misinterpreting the visual channels such as colors, sizes, or shapes. For instance, a viewer might misalign a data point with its corresponding axis value, leading to inaccurate conclusions. Similarly, misconstruing the visual channels can result in a viewer associating a color with the wrong category or mistaking the size of a visual element for a different data dimension than intended. Decoding barriers can either stem from misconceptions at earlier stages of the process, or happen in isolation as a result of misattention. Decoding barriers can fall into one of three categories:

- D1: Reads wrong data element
- D2: Reads wrong visual channel
- D3: Reads wrong value off of axis or colormap

4.5.1 ***D1: Reads wrong data element.*** Reading off the wrong data element within a visualization often traces back to earlier barriers barriers: (1) a flawed mental representation of the visualization - *E1: Misunderstanding the visual encodings* (See Figure 5a for example), (2) a misreading of axes, common in histograms or time series - *E2: Misinterprets the axis or colormap*, or (3) misunderstanding or misreading the primary query, as captured by the two translation barriers *T1: Misunderstanding the question* (See Figure 4a for example) and *T2: Misreading the question* (See Figure 4b for example). For

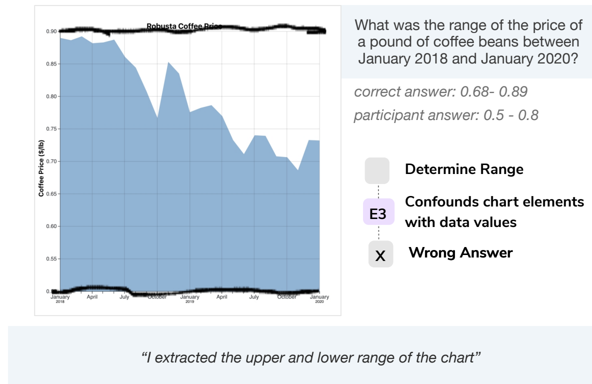

(a) Example of an *E3: Confounds chart elements with data values* mistake. This participant read the min and max values of the axis as the min and max values in the dataset.

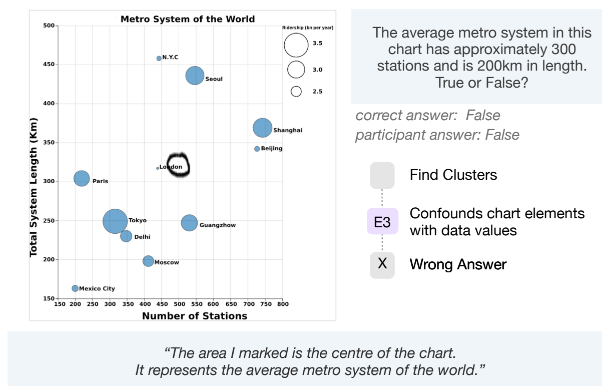

(b) Example of an *E3: Confounds chart elements with data values* mistake. This participant estimated the center of the chart as the average value for the data.

Figure 7: Examples of Encoding Barriers encountered by participants in the survey. The sketches portray participant annotations, while the pathway on the right of each panel denotes the sequence of mistakes made by the participant.

example, viewers who struggled with interpreting axes (E2 barrier) were also inclined to read the wrong data elements in time series and histograms. When the foundational conceptual error was a translation barrier, categorical visualizations, such as bar charts showcasing internet speeds across nations or stacked bars detailing room service costs in global cities, were most frequently misread.

4.5.2 ***D2: Reads wrong visual channel.*** In this barrier, viewers typically focus on the correct data element but read off the wrong visual channel. One frequent mistake is referencing the y-axis in a scatterplot when the x-axis references the correct data attribute (or vice versa), especially when both have comparable value ranges. A more pronounced instance of this type of mistake involves interpreting an entirely different visual channel, like size, when the question pertains to length or position.

This error, much like the D1 barrier where the wrong data element is read, can originate from previous conceptual misunderstandings. However, our survey's qualitative feedback suggested that attention lapses often played a significant role in these types of mistakes. Case in point, when participants were prompted to illustrate their answers on the visualization, many recognized their oversight. The act of drawing consistently made participants realize their own errors, a theme that prominently emerged in our study, which we delve deeper into in Section 6.

4.5.3 ***D3: Reads wrong value off of axis or colormap.*** The final decoding barrier emerges when viewers correctly identify the data element and read the correct visual channel but misread its value from the axis, legend, or colormap. This error often stems from an earlier encoding barrier, particularly *E2: Misinterprets the axis or colormap*. However, it can also arise from less egregious conceptual missteps, like errors in estimation.

4.6 Unraveling Misconceptions: A Qualitative Dive into Participants' Rationale

In order to better understand the types of conceptual flaws that led to poor accuracy in each of the charts, we analyzed the percentage of each misconception code, for each chart and task, as presented in Figure 8.

We first investigate the reasoning for incorrect task completions with stacked charts, which exhibited the lowest average accuracy of all charts. The most prevalent code in all stacked charts was *E1: Misunderstood the visual encoding*. This conceptual barrier was most prevalent for stacked area charts, which often resulted from participants interpreting the chart as *overlaid* areas instead of *stacked* areas. An interesting, albeit incorrect, interpretation of the stacked bar charts was that each bar represented the *range* of possible values for that category. So in the case of our dataset, *Peanuts in Seoul cost between 35 and 42 Dollars* (direct quote from one of the participants). Other participants simply read the top of the bar, the bottom of the bar, or even the middle of the bar as an 'average' value. Stacked bar charts also suffered from the operational decoding barrier *D1:Reads wrong data element*, which is when participants would look at the wrong segment or the wrong bar entirely. This is likely due to the larger amount of data (10 cities and 5 categories) and the perceptually more challenging task of isolating a given category across all bars.

Patterns from the qualitative analysis reveal unique mistakes associated with certain chart types, such as the Choropleth map, Bubble charts, and Pie charts. The Choropleth map had an average accuracy of 94% in our study. Yet, for those who made errors, the challenges centered on misinterpreting values from the quantized colormap (*E2:Misinterpreted the Axis/Colormap*). Some participants struggled with the idea that each color represented a range of values, such as 4-6%. They expressed confusion about selecting the right 'interval' from dropdown options, expecting a single value representation for each state's color.

Meanwhile, Pie charts, with an average accuracy of 90%, presented one predominant type of error: *D3:Reading wrong values off of axis*. Participants often correctly identified the relevant pie slice but misconstrued the slice's angle in degrees as its percentage value. For instance, a slice taking up a quarter of the pie (or 90 degrees) was erroneously seen as representing 90% of the sales. This mistake underscores an interesting aspect of visualization literacy: participants might latch onto familiar metrics (like degrees) but falter when bridging that understanding to another concept (like percentage distribution). This mistake was exclusive to the tasks of retrieving values. When asked to identify extremes, participants only had to find the largest/smallest slice of the pie, which did not challenge their ability to decode absolute values from the categories.

The Bubble chart is the only chart in our survey that encodes 3 data properties instead of 2 (encoded in the x/y position and size of the circle mark). The most common misconception with this chart type was *D2:Read Wrong Visual Channel*. For example, participants would prioritize the size of the circle over their position along the x or y axis. The Bubble chart was also the one with the highest percentage of this type of error among all chart types. Analyzing the types of mistakes per visualization task (right panel in Figure 8) reveals that the two tasks where viewers consistently read off the wrong visual channel in the Bubble chart were finding extremes (selecting the largest circle instead of the one on the extremities of the axis) as well as making comparisons (selecting the larger of two circles instead of the one farther to the right/top on an axis).

5 CHARTING PERFORMANCE: ACCURACY AND TIME METRICS ACROSS VISUAL REPRESENTATIONS

We processed the quantitative data from our survey at two levels of granularity: chart-centric and task-centric. We first evaluated the average accuracy and time across each chart type (Figure 9) and subsequently delved into task-specific performance within each chart (Table 2). We also performed a time-centric analysis to disambiguate the effect of barriers on accuracy and time taken to perform a task (Figure 10).

Chart Specific Performance Trends. At a high level, the chart-specific analysis reveals that participants found the stacked area chart the most challenging (average accuracy 47%), while they performed best with pie, line, and choropleth charts (average accuracies > 90%). Charts with lower accuracy scores typically required slightly more time for participants to interpret. Additionally, for charts with lower average accuracies, there was often a greater distribution of scores amongst tasks (grey markers in Figure 9). That is, participants performed a few tasks with the chart very well, and others very poorly. We present a more detailed task-centric analysis of the results in the section below.

At the highest end of the performance scale, Line charts, Pie charts, and Choropleth maps consistently exhibited a mean accuracy of over 90%. This suggests that users typically perform better when interpreting more familiar chart types, reflected in both their accuracy and efficiency. Our qualitative analysis of sketches and rationale for these charts show that mistakes often arise not from the inherent complexity of the chart or its design, but primarily from participants misinterpreting or misreading the question, leading them to extract the incorrect data element from the visualization.

The three types of charts that participants struggled with the most in terms of average accuracy were stacked area charts, area charts, and stacked bar charts. Notably, the stacked 100% bar chart

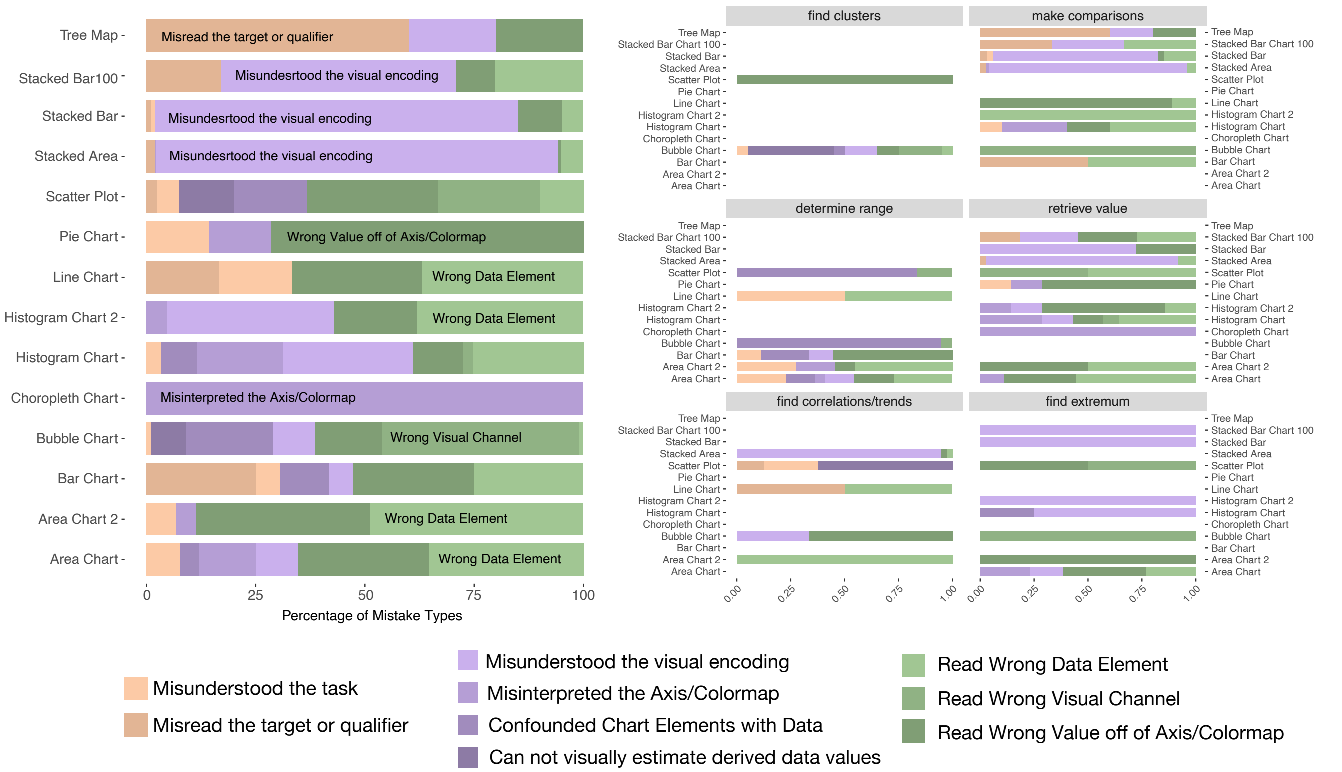

Figure 8: Distribution of Qualitative Codes for Each Chart and Task Type. This 100% stacked bar chart depicts the frequency distribution of qualitative codes identified across different chart types in the study (left panel) and for different visualization tasks (right panel). Each segment within the bars represents a distinct code, allowing for a comparative analysis of prevalent themes and misconceptions associated with each chart. In the left panel, the most prevalent code for each chart is highlighted when they account for over 30% of the misconceptions observed for that chart.

also presented significant difficulties, ranking as the fifth most challenging for participants to interpret. This suggests that visualizations that layer data, especially in a stacked format, can be particularly tricky for users to decode, potentially due to the overlapping elements and compounded data representations. Our analysis of the qualitative data confirms this and suggests that participants often do not understand the visual encodings in stacked charts. The single task on which participants performed worst was retrieving values from the stacked area chart, with an average accuracy of 10% (Table 2).

Ultimately, stacked charts, while visually appealing and effective for conveying cumulative totals, can pose significant challenges in interpretation, especially when readers aim to compare individual segments within the bar or across bars. The challenge participants encountered was trying to compare the sizes of segments that were not aligned to a common baseline. Additionally, stacked charts that displayed multiple categories (such as the stacked bars with 8 cities and 5 categories) led to an even higher incidence of mistakes.

The Effect of Barriers on Task Completion Time. We also investigated how participants who answered tasks correctly compared to those who did not in terms of time to complete the task. That is, do barriers to visual literacy tend to increase the time a participant spends on performing the task? Our results show that yes, for 10 of the 14 charts, incorrect tasks took longer than correct ones (Figure 10). Charts where the correct task completion took longer included the bar chart, the stacked bar chart 100, the bubble chart, and the stacked bar chart. The average difference in time between correct and incorrect tasks varied by chart, from less than 2 seconds for the stacked area chart tasks up to over 25 seconds for the bubble chart tasks.

One notable trend is that for the more challenging charts (lower overall accuracy, right side of Figure 10), the time difference between people who got the tasks correct and incorrect tends to be small. That is, participants who got it wrong were only a little slower than those who got it right. Conversely, in 'simpler' charts (high average accuracy, left side of Figure 10), the average difference between the two groups tends to be higher. Participants who got the tasks wrong in the simple charts took up to 20 seconds longer than those who got it right. While the absolute number of people who got tasks wrong for simple charts is much smaller (size of red circles in Figure 10), those who did took much longer to answer.

There are a few exceptions to this trend, mostly within the more challenging charts, where the slower participants were those who answered correctly. A particularly egregious example is the bubble

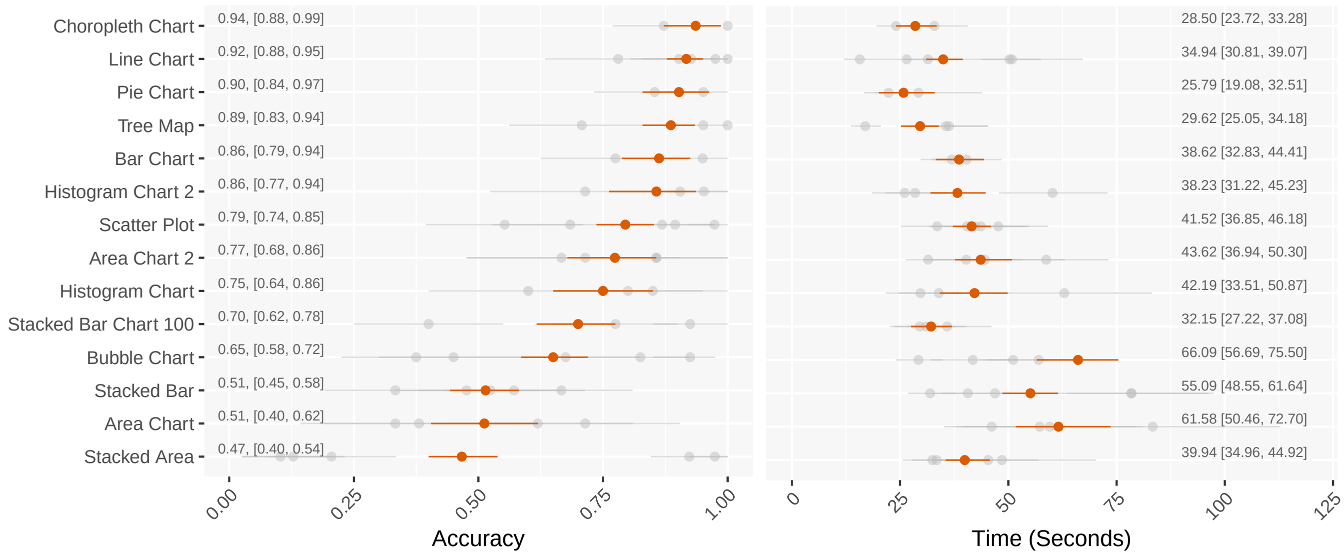

Figure 9: Task Performance Analysis by Chart. The left panel illustrates mean task accuracy, while the right showcases the time taken. Red circles denote overall mean accuracy across all tasks, while grey circles represent mean accuracy for individual tasks within each chart. Bars indicate a 95% confidence interval. Chart types are organized based on their overall mean accuracy.

chart, where participants who got the tasks correct took an average of 25 seconds longer. A closer analysis of this data broken down by task reveals that this difference is mostly driven by the task where participants were asked to compare an attribute for two cities (BuCQ3, Table 2). To be performed correctly, this question required close attention to the fact that the attribute in question was encoded in size, and not in either axis. Analysis of barriers for the Bubble Chart shows that *D2:Reads the wrong visual channel* was the most common barrier encountered when participants completed the task incorrectly (Figure 8). We hypothesize that participants who answered the question correctly took the time to analyze all visual channels and decode the correct one. Two other, less stark, examples of this same pattern are the two stacked bar charts (standard and 100%), where participants who performed the tasks correctly took on average 5 and 9 seconds longer than those who did not, respectively. Given the complexity and oftentimes novelty of stacked bar charts, we hypothesize that participants who took the time to understand the encoding tended to take a few seconds longer and complete the task accurately.

Task-Centric Analysis of Performance. We also adopted an orthogonal analysis method, categorizing questions based on the visualization tasks they encompassed, such as value retrieval and comparison. Consider the task of identifying clusters, presented to participants using both a scatterplot (with a mean accuracy of 0.87) and a bubble chart (with a mean accuracy of 0.38). The pronounced disparity in performance between the scatter plot and the bubble chart implies that the additional visual channel of size in the bubble chart can be a source of confusion. This observation was supported in the qualitative data, where some participants ascribed exaggerated significance to larger circles, skewing their interpretation.

Another interesting juxtaposition in performance for the same task relates to making comparisons. The stacked area chart had the poorest performance for this task, registering an average accuracy of 13% and 21% for the two tasks of this type. This is likely attributed to the inherent challenge of comparing data points with varied baselines. In contrast, the stacked bar chart, which poses a similar encoding challenge, performed relatively better, achieving an average accuracy of 48% and 67% in the comparison tasks. The qualitative feedback from participants suggests that the continuous presentation of data in the stacked area chart led many viewers to mistakenly assume that all categories started from a baseline of 0, rather than the peak of the preceding category.

The quantitative data from the survey provided a foundational understanding of which chart types and tasks posed greater difficulties for participants. Meanwhile, the qualitative analysis of participants' sketches and textual feedback, discussed in the previous section, provided rich insight into the underlying cognitive barriers and misconceptions leading to these performance disparities.

Participant-Centric Analysis. Among the demographic information collected from participants was their self-reported level of familiarity with visualizations, which ranged from 'no familiarity' to 'expert'. The number of participants in each group can be seen in Table 1. We conducted an analysis of overall participant performance on the VLAT as a function of their familiarity with visualizations as well as the frequency of each type of barrier in the different groups (Table 1). From an accuracy standpoint, average VLAT scores increased with increasing visualization literacy. More specifically, relatively skilled and expert participants scored an average of 80% compared to 68% for participants with no familiarity with visualizations. In the analysis of barriers for each

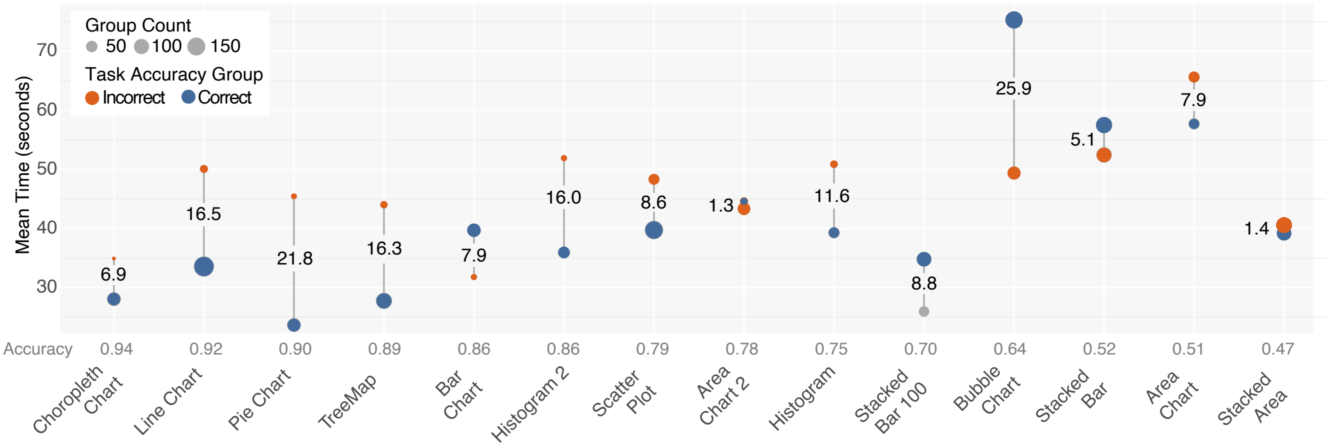

Figure 10: Differences in time for correct (blue) and incorrect (red) task completions. Average time taken per chart type, grouped by correct and incorrect task completions. The distance between the two markers for each chart captures how much faster one group was than the other. Circles are sized by the number of task completions in that group. Labels indicate the difference in time between the two groups. Chart types are organized based on their overall mean accuracy

group, the frequency of translation errors (orange bars in Table 1 decreased steadily as familiarity with visualizations increased. This suggests that the ability to translate an analytic task to a visual one is skill that is inherent to increased visualization literacy. Another noteworthy pattern in the data was that for all groups (except for expert, which had only 1 participant), the most common type of barrier encountered was E1:Misunderstood the visual encoding. This finding accentuates the need to weave visualization education into mainstream curricula and broaden the general public's familiarity with diverse chart formats. We discuss this in more detail in the following section.

6 DISCUSSION

As we chart the challenges to visualization literacy, we find that several barriers can impede accurate interpretation and understanding of the underlying data. Here, we reflect on these barriers, their intersection with related efforts in visualization, and implications for future work in the field of visualization education.

Intersection of Barriers in Existing Work. This work outlined the three main types of barriers that can occur when a viewer extracts incorrect or no information from a visualization. Existing work in the field of visualization has often encountered one or more of these barriers, to varying degrees of granularity, scope, and rigor. For example, Grammel et al. [21] investigated how novices build visualizations and found three major barriers faced by the participants: (1) translating the task into data attributes (similar to our translation barrier), designing visual mappings (closely related to encoding barriers when reading a visualization) and interpreting visualizations (a higher level category that encapsulates our decoding barriers). Participants in their study were tasked to *create* visualizations and, therefore, had to employ a different set of skills than when *reading* a visualization. Nevertheless, the overlap in the types of barriers encountered suggests that education efforts that address viewers' ability to translate, encode, and decode are a valuable contribution to the field.

Other efforts investigated barriers to visualization for specific visualization types ([2, 17, 29]). For parallel coordinates, for example, Peebles et al. [29] found two main barriers: (1) a lack of understanding of coordinate systems and (2) visual complications from crossing lines and clutter. We can draw a parallel between their first barrier and our concept of encoding barriers, which encompass difficulties in understanding the layout. Their second barrier is particularly interesting in that it provides a different precursor to decoding barriers than the ones observed in our study. The decoding errors observed in our study mostly stemmed from either inattention or a prior misconception of how the data was encoded. The issue of visual clutter identified by Peebles et al. serves as an example of the broader category of perceptual aspects, which can often lead to decoding errors. Understanding the factors, whether in the visualization or in the viewer, that lead to the different barriers is critical in providing the appropriate educational intervention.

In investigating what triggered specific barriers in our dataset, we observed that participants are often susceptible to more than one barrier, with a first misconception often triggering a second or third barrier during analysis. More specifically, we observe that mistakes early in the visual analysis process, such as misinterpreting colors, sizes, or the layout, often cascade into subsequent stages of data interpretation, such as decoding data values. Educational interventions can address many of the conceptual barriers we have identified at their root.

Mitigating Barriers through Visualization Education. In a recent overview of challenges and opportunities in data visualization education, Bach et al. [5] highlight the need to improve visualization literacy through different educational approaches. Our work can provide guidance on the types of barriers that these efforts can target. For example, educational approaches that target encoding barriers can provide guidance for viewers on how to interpret a

Visualization Familiarity	Average VLAT Score	Frequency of Visualization Barriers
No familiarity (5)	0.680	
Beginner (35)	0.693	
Moderately knowledgeable (56)	0.709	
Relatively Skilled (25)	0.803	
Expert (1)	0.800	

Table 1: Table summarizing performance metrics and the occurence of visualization barriers across different self-reported familiarity with visualization. For each group, the number in parenthesis indicates how many participants self-reported into that category. Frequencies of barriers span 0 to 100%.

visual channel. Decoding barriers might be minimized by the effective use of annotations. A more active approach to address barriers might walk a viewer from the data to visual encoding in a simplified example, followed by exercises for the user to gain hands-on experience.

A promising avenue for learning, both in the field of visualization and elsewhere, is active learning [34]. By allowing users to engage directly with visual elements and receive immediate feedback, it becomes possible to correct misconceptions in real-time. Active learning has been successfully applied in targeted efforts at enhancing visualization literacy across different settings and for different target ages [6, 12, 23]. Interactivity, especially in the realm of visualization literacy, provides users with a tangible and immediate connection to abstract data, enabling a more hands-on learning experience. This type of learning can be particularly useful for addressing encoding barriers, as it helps viewers better understand the underlying relationship between the data and the visual channels. Moreover, as they manipulate the visual elements, users can validate or challenge their initial understandings, allowing for a continuous learning process. Integrating tutorials or prompts that guide users in this exploration can further streamline the learning curve, ensuring that they not only view data but also understand the underlying patterns and narratives. As visualization tools evolve, the seamless incorporation of these interactive features will undoubtedly play a pivotal role in bridging the gap between complex data and its interpretation.

A particular type of active learning that has proven useful in the field of education is sketching [1, 24]. Interestingly, the qualitative analysis of the survey results revealed a pattern around participants' self-realization of errors when they were prompted to sketch directly on the visualization. Specifically, in 39 instances (i.e., 8% of incorrect responses), participants initially provided an incorrect answer. However, when asked to elucidate their thinking process by drawing on the visualization, they identified the error in their initial assessment. These instances happened for both operational decoding barriers (i.e., reading the wrong data value) as well as for conceptual encoding barriers, where drawing on the visualization helped them understand the encoding. For example, one participant said: *"When I drew the lines, I noticed that I answered the wrong month. So l realized that I was confused by the way the months were organized"*. This result aligns with earlier research that has explored the value of sketching as a way of prompting reflection [1, 24]. Our results underscore the pivotal role of sketching in pinpointing conceptual missteps.

Ultimately, our study outlines the multifaceted nature of barriers in visualization literacy, and we emphasize the critical role of targeted education in overcoming these challenges. Our findings provide support for an educational approach that combines theoretical understanding with practical, hands-on experiences like active learning and sketching. These methods not only aid in clarifying concepts but also foster a deeper engagement with the material, allowing learners to internalize and apply their knowledge more effectively.

Visualization	Task ID	Vis Task	Time (s)	Acc (%)	Accuracy Summary
Area chart	ACQ1	retrieve value	46	71	
	ACQ2	find extremum	60	62	
	ACQ3	determine range	83	38	
	ACQ4	find correlations/trends	57	33	
Area chart 2	AC2Q4	find correlations/trends	45	86	
	AC2Q2	find extremum	31	86	
	AC2Q1	retrieve value	40	71	
	AC2Q3	determine range	59	67	
Bar chart	BCQ4	make comparisons	37	95	
	BCQ3	determine range	40	78	
Bubble Chart	BuCQ3	make comparisons	158	93	
	BuCQ6	find extremum	29	82	
	BuCQ4	find correlations/trends	42	68	
	BuCQ2	determine range	51	45	
	BuCQ5	find clusters	57	38	
Choropleth chart	CCQ1	find extremum	24	100	
	CCQ2	retrieve value	33	87	
Histogram Chart	HCQ3	find extremum	30	95	
	HCQ2	make comparisons	34	90	
	HCQ1	retrieve value	63	71	
Histogram Chart 2	HC2Q2	make comparisons	28	85	
	HC2Q3	find extremum	26	80	
	HC2Q1	retrieve value	60	60	
Line chart	LCQ2	find extremum	16	100	
	LCQ1	retrieve value	27	98	
	LCQ4	find correlations/trends	31	93	
	LCQ3	determine range	51	90	
	LCQ5	make comparisons	50	78	
Pie chart	PCQ2	find extremum	22	95	
	PCQ1	retrieve value	29	85	
Scatter Plot	SPQ1	retrieve value	34	97	
	SPQ2	find extremum	42	89	
	SPQ4	find clusters	44	87	
	SPQ3	determine range	48	68	
	SPQ5	find correlations/trends	40	55	
Stacked area	SAQ4	find correlations/trends	49	97	
	SAQ3	find extremum	32	92	
	SAQ2	make comparisons	45	21	
	SAQ5	make comparisons	33	13	
	SAQ1	retrieve value	40	10	
Stacked Bar	SBQ4	make comparisons	78	67	
	SBQ2	find extremum	32	57	
	SBQ1	retrieve value	78	52	
	SBQ3	make comparisons	41	48	
	SBQ5	retrieve value	47	33	
Stacked Bar 100	SB100Q2	make comparisons	30	93	
	SB100Q1	retrieve value	36	78	
	SB100Q3	find extremum	31	40	
TreeMap	TMQ3	identify hierarchical structure	17	100	
	TMQ1	find extremum	36	95	
	TMQ2	make comparisons	36	71	

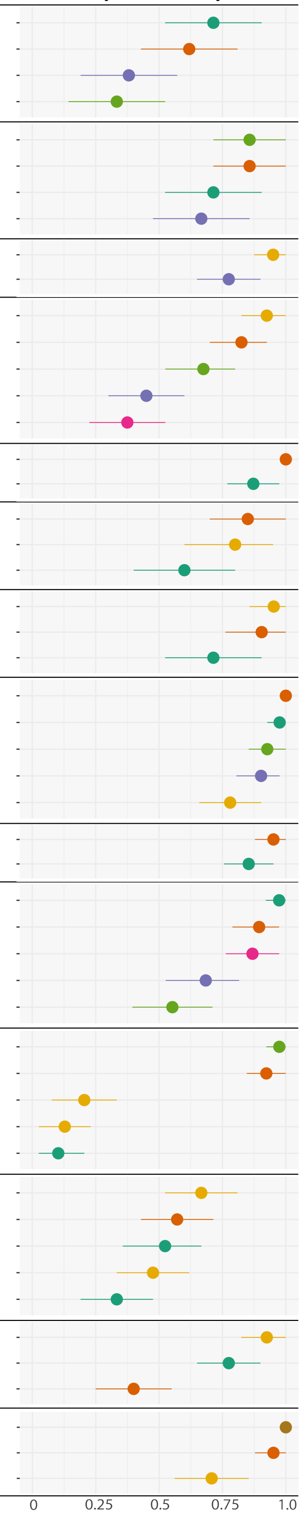

Table 2: Table summarizing performance metrics across different chart types. Columns detail the chart name, Task ID, specific visualization task addressed, average response time in seconds, average accuracy as a percentage, and an embedded chart visualizing accuracy rates for each chart type, color-coded by visualization task. For each chart, tasks are sorted from highest to lowest accuracy

7 CONCLUSIONS AND FUTURE WORK

In this study, we conducted an in-depth exploration of the barriers to visualization literacy, with an emphasis on understanding the underlying conceptual challenges. Drawing from the extensive qualitative data collected in our survey, our findings shed light on the specific misconceptions and difficulties individuals face across various visualization types. Our work further delves into the unique barriers associated with 12 different chart types, providing an overview of the nuances in comprehension. Through our research, we have developed a classification system that outlines the conceptual and operational barriers individuals encounter throughout the visualization interpretation process. A complementary finding was the role of sketching in fostering visualization literacy, exemplified by instances where participants' self-correction was prompted by their own sketches. Our findings not only contribute to the existing body of knowledge but also suggest potential avenues for improving visualization design and education.

Future research should continue to explore these barriers across diverse contexts, populations, and visualization types. We see our study as an initial step, paving the way for more in-depth research aimed at enhancing visualization literacy. Our classification, we believe, lays the groundwork for developing more effective methods to teach visualization literacy. A comprehensive understanding of the prevalent barriers and their underlying causes will enable the creation of more targeted interactive tutorials that address these common misconceptions. Looking ahead, there is potential for the development of dynamic systems that can proactively assist by detecting user misconceptions through their interactions.

In this work, we use specific analytic tasks, as captured by the VLAT, to investigate barriers to understanding. We take this approach to more easily isolate instances of misunderstandings by tracking incorrect task completions. However, real-world analysis can also have a more exploratory nature, where viewers are not trying to perform a specific task. Future work can explore the types of barriers that viewers can face in these settings as well, and how they overlap with the ones identified in this study. Additionally, to reduce survey fatigue for our participants, we designed the study to only ask for rationale and sketches from participants who perform a task incorrectly. However, future versions of this work that also capture sketches and rationale from correct task completions can provide rich insight into different strategies both within and across participants for correct answers. Future surveys can also analyze how participants perform on tasks with difference levels of difficulty and discrimination, as described by CTT (Classical Test Theory) and IRT (Item Response Theory).

An additional avenue for future work is eye-tracking studies, which can provide valuable insight for understanding mental models in visualization literacy. Related work in this area has looked at how eye movements can reveal cognitive processes [38] as well as uncovered the impact of design choices in visualizations, such as the choice of using linear or log scales [20]. By monitoring where and for how long individuals gaze at particular elements of a visualization, future work can gain insight into the cognitive processes and strategies users employ when interpreting visual data. Such studies can reveal the sequence in which information is processed, areas of confusion or misinterpretation, and the elements that capture users' attention the most. When integrated with other research methods, eye-tracking can provide a detailed picture of the underlying mental models that drive users' interactions with visualizations and shed light on areas where intervention or redesign might be beneficial.

As we navigate an age replete with data, ensuring that individuals can accurately interpret visual representations is crucial. By understanding and addressing the barriers to visualization literacy, we pave the way for more informed decisions and a deeper appreciation of the stories that data can tell.

REFERENCES

[1] Gregor Aisch, Amanda Cox, and Kevin Quealy. 2015. You Draw It: How Family Income Predicts Children's College Chances. *The New York Times* 28 (2015).
[2] Mashael AlKadi, Vanessa Serrano, James Scott-Brown, Catherine Plaisant, Jean-Daniel Fekete, Uta Hinrichs, and Benjamin Bach. 2022. Understanding Barriers to Network Exploration With Visualization: A Report from the Trenches. *IEEE Transactions on Visualization and Computer Graphics* 29, 1 (2022), 1–11. https://doi.org/10.1109/TVCG.2022.3209487
[3] Basak Alper, Nathalie Henry Riche, Fanny Chevalier, Jeremy Boy, and Metin Sezgin. 2017. Visualization Literacy at Elementary School. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, Denver Colorado USA, 5485–5497. https://doi.org/10.1145/3025453.3025877
[4] Lynna J. Ausburn and Floyd B. Ausburn. 1978. Visual Literacy: Background, Theory and Practice. *Programmed Learning and Educational Technology* 15, 4 (Nov. 1978), 291–297. https://doi.org/10.1080/0033039780150405
[5] Benjamin Bach, Mandy Keck, Fateme Rajabiyazdi, Tatiana Losev, Isabel Meirelles, Jason Dykes, Robert S. Laramee, Mashael AlKadi, Christina Stoiber, Samuel Huron, Charles Perin, Luiz Morais, Wolfgang Aigner, Doris Kosminsky, Magdalena Boucher, Søren Knudsen, Areti Manataki, Jan Aerts, Uta Hinrichs, Jonathan C. Roberts, and Sheelagh Carpendale. 2023. Challenges and Opportunities in Data Visualization Education: A Call to Action. arXiv:2308.07703 [cs]
[6] S. Sandra Bae, Rishi Vanukuru, Ruhan Yang, Peter Gyory, Ran Zhou, Ellen Yi-Luen Do, and Danielle Albers Szafir. 2022. Cultivating Visualization Literacy for Children Through Curiosity and Play. https://doi.org/10.48550/arXiv.2208.05015 arXiv:2208.05015 [cs]
[7] Ryan Birchfield, Maddison Caten, Errica Cheng, Madyson Kelly, Truman Larson, Hoan Phan Pham, Yiren Ding, Noelle Rakotondravony, and Lane Harrison. 2022. VisQuiz: Exploring Feedback Mechanisms to Improve Graphical Perception. In *2022 IEEE Visualization and Visual Analytics (VIS)*. IEEE, Oklahoma City, OK, USA, 95–99. https://doi.org/10.1109/VIS54862.2022.00028
[8] Katy Börner, Adam Maltese, Russell Nelson Balliet, and Joe Heimlich. 2016. Investigating Aspects of Data Visualization Literacy Using 20 Information Visualizations and 273 Science Museum Visitors. *Information Visualization* 15, 3 (July 2016), 198–213. https://doi.org/10.1177/1473871615594652
[9] Jeremy Boy. 2015. *Engaging the People to Look Beyond the Surface of Online Information Visualizations of Open Data*. Telecom ParisTech.
[10] Jeremy Boy, Ronald A. Rensink, Enrico Bertini, and Jean-Daniel Fekete. 2014. A Principled Way of Assessing Visualization Literacy. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 1963–1972. https://doi.org/10.1109/TVCG.2014.2346984
[11] EVA BRUMBERGER. 2004. The Rhetoric of Typography: Effects on Reading Time, Reading Comprehension, and Perceptions of Ethos. *Technical Communication* 51, 1 (2004), 13–24. jstor:43089069
[12] Vetria Byrd and Nicole Dwenger. 2021. Activity Worksheets for Teaching and Learning Data Visualization. *IEEE Computer Graphics and Applications* 41, 6 (Nov. 2021), 25–36. https://doi.org/10.1109/MCG.2021.3115396
[13] Bum chul Kwon, Brian Fisher, and Ji Soo Yi. 2011. Visual Analytic Roadblocks for Novice Investigators. In *2011 IEEE Conference on Visual Analytics Science and Technology (VAST)*. Institute of Electrical and Electronics Engineers (IEEE), Providence, RI, USA, 3–11. https://doi.org/10.1109/VAST.2011.6102435
[14] Juliet Corbin. 1990. *Basics of Qualitative Research Grounded Theory Procedures and Techniques*. Newbury Park, California Sage.
[15] Yuan Cui, Lily W. Ge, Yiren Ding, Fumeng Yang, Lane Harrison, and Matthew Kay. 2023. Adaptive Assessment of Visualization Literacy. arXiv:2308.14147 [cs]
[16] Peter Felten. 2008. Visual Literacy. *Change: The Magazine of Higher Learning* 40, 6 (Nov. 2008), 60–64. https://doi.org/10.3200/CHNG.40.6.60-64
[17] Elif E. Firat, Alena Denisova, and Robert Laramee. 2020. *Treemap Literacy: A Classroom-Based Investigation*.
[18] Elif E Firat, Alark Joshi, and Robert S Laramee. 2022. Interactive Visualization Literacy: The State-of-the-Art. *Information Visualization* 21, 3 (July 2022), 285–310. https://doi.org/10.1177/14738716221081831

[19] Elif E. Firat, Alark Joshi, and Robert S. Laramee. 2022. VisLitE: Visualization Literacy and Evaluation. *IEEE Computer Graphics and Applications* 42, 3 (May 2022), 99–107. https://doi.org/10.1109/MCG.2022.3161767
[20] Joseph Goldberg and Jonathan Helfman. 2011. Eye Tracking for Visualization Evaluation: Reading Values on Linear versus Radial Graphs. *Information Visualization* 10, 3 (July 2011), 182–195. https://doi.org/10.1177/1473871611406623
[21] Lars Grammel, Melanie Tory, and Margaret-Anne Storey. 2010. How Information Visualization Novices Construct Visualizations. *IEEE Transactions on Visualization and Computer Graphics* 16, 6 (Nov. 2010), 943–952. https://doi.org/10.1109/TVCG.2010.164
[22] Hacer Hanci. 2022. Investigation of High School Students' Visual Literacy Levels. *International Journal of Research in Education and Science* 8, 3 (2022), 611–625.
[23] Samuel Huron, Benjamin Bach, Uta Hinrichs, Mandy Keck, and Jonathan C Roberts. 2020. IEEE VIS Workshop on Data Vis Activities to Facilitate Learning, Reflecting, Discussing, and Designing. (2020).
[24] Yea-Seul Kim, Katharina Reinecke, and Jessica Hullman. 2017. Explaining the Gap: Visualizing One's Predictions Improves Recall and Comprehension of Data. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, Denver Colorado USA, 1375–1386. https://doi.org/10.1145/3025453.3025592
[25] Sukwon Lee, Sung-Hee Kim, Ya-Hsin Hung, Heidi Lam, Youn-Ah Kang, and Ji Soo Yi. 2016. How Do People Make Sense of Unfamiliar Visualizations?: A Grounded Model of Novice's Information Visualization Sensemaking. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (Jan. 2016), 499–508. https://doi.org/10.1109/TVCG.2015.2467195
[26] Sukwon Lee, Sung-Hee Kim, and Bum Chul Kwon. 2017. VLAT: Development of a Visualization Literacy Assessment Test. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (Jan. 2017), 551–560. https://doi.org/10.1109/TVCG.2016.2598920
[27] Sukwon Lee, Bum Kwon, Jiming Yang, Byung Lee, and Sung-Hee Kim. 2019. The Correlation between Users' Cognitive Characteristics and Visualization Literacy. *Applied Sciences* 9, 3 (Jan. 2019), 488. https://doi.org/10.3390/app9030488
[28] Saugat Pandey and Alvitta Ottley. 2023. Mini-VLAT: A Short and Effective Measure of Visualization Literacy. *Computer Graphics Forum* 42, 3 (June 2023), 1–11. https://doi.org/10.1111/cgf.14809
[29] D. Peebles, D. Ramduny-Ellis, G. Ellis, and J. V. H. Bonner. 2013. *The Influence of Graph Schemas on the Interpretation of Unfamiliar Diagrams.* Technical Report. UK: School of Human and Health Sciences, University of Huddersfield.
[30] Marc Satkowski, Franziska Kessler, Susanne Narciss, and Raimund Dachselt. 2022. Who Benefits from Visualization Adaptations? Towards a Better Understanding of the Influence of Visualization Literacy. In *2022 IEEE Visualization and Visual Analytics (VIS)*. IEEE, Oklahoma City, OK, USA, 90–94. https://doi.org/10.1109/VIS54862.2022.00027
[31] R.C. Schank. 1994. Active Learning through Multimedia. *IEEE MultiMedia* 1, 1 (1994), 69–78. https://doi.org/10.1109/93.295270
[32] Dana Statton Thompson, Stephanie Beene, Katie Greer, Mary Wegmann, Millicent Fullmer, Maggie Murphy, Sara Schumacher, and Tiffany Saulter. 2022. A Proliferation of Images: Trends, Obstacles, and Opportunities for Visual Literacy. *Journal of Visual Literacy* 41, 2 (April 2022), 113–131. https://doi.org/10.1080/1051144X.2022.2053819
[33] Suzanne Stokes. 2002. Visual Literacy in Teaching and Learning: A Literature Perspective. *Electronic Journal for the integration of Technology in Education* (2002), 19.
[34] Yuzuru Tanahashi, Nick Leaf, and Kwan-Liu Ma. 2016. A Study On Designing Effective Introductory Materials for Information Visualization. *Computer Graphics Forum* 35, 7 (2016), 117–126. https://doi.org/10.1111/cgf.13009
[35] Conrad Taylor. 2003. New Kinds of Literacy, and the World of Visual Information. *EIGVIL workshop* (2003).
[36] J.J. Van Wijk. 2006. Views on Visualization. *IEEE Transactions on Visualization and Computer Graphics* 12, 4 (July 2006), 421–432. https://doi.org/10.1109/TVCG.2006.80
[37] Leni Yang, Cindy Xiong, Jason K. Wong, Aoyu Wu, and Huamin Qu. 2023. Explaining with Examples: Lessons Learned from Crowdsourced Introductory Description of Information Visualizations. *IEEE Transactions on Visualization and Computer Graphics* 29, 3 (March 2023), 1638–1650. https://doi.org/10.1109/TVCG.2021.3128157
[38] Alfred L. Yarbus. 1967. Eye Movements During Perception of Complex Objects. In *Eye Movements and Vision*, Alfred L. Yarbus (Ed.). Springer US, Boston, MA, 171–211. https://doi.org/10.1007/978-1-4899-5379-7_8