Eliciting High-Level Visual Comprehension: A Qualitative Study

Ghulam Jilani Quadri*
University of North Carolina at Chapel Hill

Danielle Albers Szafir†
University of North Carolina at Chapel Hill

Abstract

A visualization designer creates a given visualization with a specific analytical or communication goal. However, perceptual studies of visualization effectiveness focus on isolated, low-level tasks, such as estimating specific statistics. Instead, we explore data interpretation and communication more holistically to bridge the gap between visualization designers and consumers. In this work, we conducted a qualitative study on five selected graphs from New York Times to investigate the high-level patterns people naturally see when they encounter a visualization without a guiding task. Participants described each of the tested graphs using natural language. The descriptions were coded using axial coding to identify whether the patterns people observed in the visualizations aligned with the designer's intentions. We found that interpretation varies with the number of subgraphs, additional annotations, labels, and units. Our findings provide a new lens on how design influences the high-level patterns people naturally see in visualization. The subsequent findings and guidelines on visual design can significantly strengthen social trust information visualization when data are visualized in social policies and information contexts.

Index Terms: Human-centered computing—Visualization—Visualization design and evaluation methods

1 Introduction

Information visualization helps users extract meaningful analytical insights, such as understanding and making sense of public data such as COVID-19, wildfires, or hurricanes. Visualization interpretation varies with the viewer's expertise, area of the profession, and level of education in addition to visualization design [3]. Perceptual experiments measure visual design effectiveness [5]. However, perceptual studies on the evaluation of visualization effectiveness are more focused on *low-level tasks*, typically measuring people's abilities to estimate individual, pre-specified statistical quantities. *High-level* comprehension describes the overall knowledge a viewer intuitively gains about the data without explicit cueing or guidance.

In practice, visualizations aim to provide a high-level understanding of data: designers craft visualizations to deliver a specific message about the data at a glance. However, the information people extract from a graph depends on a visualization's design. Understanding what high-level knowledge people extract from visualized information helps confirm whether the visualization designer's objective is reflected in the user's interpretation. We offer preliminary steps towards such understanding, with a vision of exploring data interpretation and communication more holistically to bridge the gap between visualization designers and consumers.

In this work, we conducted an experiment to investigate what users intuitively see in data visualization using five selected graphs (see Table 2) from New York Times [1]. We asked participants to describe what they saw in the given graphs. We recorded their verbal responses to evaluate whether the patterns participants saw matched the designer's goal as stated in the constituent articles. We found high-level understanding of visualizations depends not only on selected graphs and visual encodings but additional factors such as whether visualized data is pre-processed or statistically aggregated or whether people are familiar with the given graphs(see Sect. 3.1).

Our study's findings on the visual design can significantly strengthen social trust in information visualization when data are visualized in social policy and information contexts by improving the visualization design.

*e-mail: jiquad@cs.unc.edu
†e-mail: dnszafir@cs.unc.edu

Table 1: Summary of participants' backgrounds in the study.

Participant	Role	Occupation
P1	Expert	Professor
P2	Expert	Data Scientist
P3	Visualization User	PhD Student
P4	Visualization User	PhD Student
P5	Visualization User	PhD Student
P6	Visualization User	Graduate Student
P7	Visualization User	Graduate Student
P8	Other Participants	Software Engineer
P9	Other Participants	Software Engineer
P10	Other Participants	Software Engineer

Table 2: Information on the five graphs from New York Times [1] used in the experiment. Links to articles and graphs are embedded in the description texts.

ID	Visualization	Description	Visual Encoding
V1	Map	Covid-19 1 million death	Color saturation
V2	Area chart	Covid-19 death per wave	Area
V3	Scatterplot	Covid-19 deaths by income	Position, Color
V4	Scatterplot	Reading time for Popular text	Position
V5	Line graph	Use to tobacco products	Color, stroke or shading

2 Methodology

We selected five different visualizations—map, area chart, scatterplot (single and multiple), and a line graph, as listed in Table 2—from New York Times (see examples in Fig. 1). We recruited 10 participants (six men, four women; 21-44 years of age) with varying levels of familiarity with visualization as shown in Table 1, two of which were experts in visualization research and eight of which were more casual visualization users. We presented the sampled visualizations in random sequential order and asked participants to *"describe what you see in the graph"*. The entire study took no more than 10 minutes for each participant. The responses are recorded in audio and later coded for analysis using axial coding by two coders.

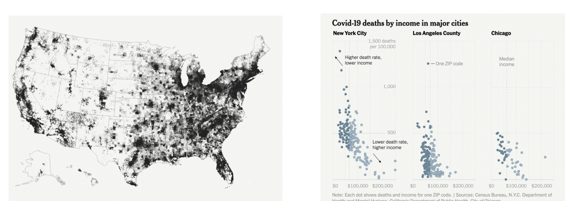

A: One million Covid-19 deaths

B: Covid-19 deaths by income

Figure 1: Examples from New York Times [1]. Graph A shows USA map plotting one million Covid-19 deaths. Graph B demonstrates the Covid-19 deaths by income.

3 RESULTS

We discuss the results of two graphs from Fig. 1 categorizing participant responses as 1) understating the graphs and 2) comments on the design, visualization, or analysis. A summary of all participant's responses is discussed at the end.

V1: The majority of participants perceived density-based information from the map. However, participants did not identify the primary objective of the map: to communicate that the crisis began in cities and spread to rural areas. Participants generally saw the COVID map as indicating regions of high and low density. Three participants (experts and a PhD student) described the map as pointing to the highly-concentrated region (see Fig. 1-A), four noted a skew of data towards the east coast, and the remaining 3 participants instead noted that the east and extreme west coast have a greater quantity or higher density of data. Four participants pointed out that the dark region represents higher quantity whereas the lighter region represents lesser quantity.

The expert participants talked in about detail how the design is misleading. For example, they mentioned Fig. 1-A lacks a scale to further interpret the approximate number of deaths in a given area. Five participants mentioned the lack of scale in the graph, and two participants asked for more information on the context of graph and data. Six participants felt this graph missed conveying more information regarding quantity and scale.

V3: The designer's objective was to show that income is a predictor of Covid-19 mortality by showing a correlation between income and both death and vaccination rates in major cities. In line with this goal, six people described the graph as showing that Covid-19 deaths are more concentrated in lower-income regions of three cities (see Sect. 2), while one participant focused on correlations between death and income from NYC alone. Two participants compared the overall number of deaths across cities. Only two participants correctly interpreted data points as representing the death and income for a given zip code, with four participants instead seeing dot color as representing incomes as being either above or below the median.

For Fig. 1, the responses varied among users when reflecting on the utility of the visualization. People felt the graphs were overloaded (three participants), complicated (two participants), the missing legend on colors (two participants), and were confused by the fact that each graph annotates different information (one participant).

Overall Alignment: We observed three types of participant responses: No match with the designer's objective (V4), partial match with the designer's objective (V1, V2, V3), and majority match with the designer's objective (V5). 90% of participants' responses for V1 and V2 partially matched the designer's intended goal. 70% of responses matched for V3. None of the participants' responses for V4 matched the designer's intended goal, whereas all responses for V5 matched the designer's intended goal.

3.1 General Discussion

The patterns people noted in the visualizations depended on the visualization's design. Visualizations presenting the data with no additional statistical or written annotations (such as a density map in Fig. 1) led to more varied participant responses. Participants generally noted that visualizations with more annotation or legends led to more consistent comprehension. As shown in Sect. 2, 3-scatterplot graphs add complexity, demanding more annotation, labels, and sound choice of visual design.

Participants' data interpretation and comprehension varied with the visualization type used. For example, in the case of a simple scatterplot, most of the participants' responses matched the designer's objective. This pattern may be because of the level of familiarity with a given graph. For example, participants' observations aligned with the intended goals more often with scatterplots (8 of 10) than with area charts (4 of 10). Similarly, people's descriptions were more aligned with designer intentions in line graphs (10 of 10) than in monochrome density maps (3 of 10).

People were less likely to observe the intended patterns in more complex graphs (see Fig. 1-B) that contained more than one graph, additional labels, and explicit statistical information. We observed that additional complexity was helpful to some participants, whereas others used negative terms like "complicated" for such graphs.

4 FUTURE WORK

Information visualization is more than visual representation: effective visualization considers the type of task [2, 4] and problem domain [4] to drives a particular design choice. One of two combinations mentioned in Schulz et al. [6]—1) *Data + Visualization = Task (?)* —points to what type of task (e.g., low-level or high-level) would be more correctly performed for a given visualization. Varied visualization design choices in Data + Visualization would decide how easily a user can understand and make sense of their data. Our results provide preliminary data about how the combination of *Data + Visualization* with the designer's objective affects the viewer's data interpretation and communication. These findings bridge the gap between visualization designers and consumers by investigating how design might intuitively support their end analytical or communicative goals. The results can guide best practices to significantly strengthen social trust in information visualization when data are visualized in social policy and information contexts by better aligning visualizations with the information they most readily communicate.

This experiment is a subset of a more extensive proposed study on more controlled stimuli; hence, we cannot generalize the findings with smaller participants. We believe a more extensive study on diverse 1) visualizations with additional design choices, 2) datasets, 3) populations, and 4) response types (text, think-aloud, eye-tracking) will allow us to more broadly compare and contrast user understanding and designer goals.

ACKNOWLEDGMENTS

This work was partially supported by the National Science Foundation under grants no. 2127309 to the Computing Research Association for the CIFellows project and IIS-2046725.

REFERENCES

[1] https://www.nytimes.com/, 2022.

[2] R. Amar, J. Eagan, and J. Stasko. Low-level components of analytic activity in information visualization. In *IEEE Symposium on Information Visualization, 2005. INFOVIS 2005.*, pp. 111–117. IEEE, 2005.

[3] K. W. Hall, A. Kouroupis, A. Bezerianos, D. A. Szafir, and C. Collins. Professional differences: A comparative study of visualization task performance and spatial ability across disciplines. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):654–664, 2021.

[4] T. Munzner. A nested model for visualization design and validation. *IEEE transactions on visualization and computer graphics*, 15(6):921–928, 2009.

[5] G. J. Quadri and P. Rosen. A survey of perception-based visualization studies by task. *IEEE Transactions on Visualization and Computer Graphics*, 2021.

[6] H.-J. Schulz, T. Nocke, M. Heitzler, and H. Schumann. A design space of visualization tasks. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2366–2375, 2013.