Why Change My Design: Explaining Poorly Constructed Visualization Designs with Explorable Explanations

Leo Yu-Ho Lo ⓘD, Yifan Cao ⓘD, Leni Yang ⓘD, and Huamin Qu ⓘD

Abstract—Although visualization tools are widely available and accessible, not everyone knows the best practices and guidelines for creating accurate and honest visual representations of data. Numerous books and articles have been written to expose the misleading potential of poorly constructed charts and teach people how to avoid being deceived by them or making their own mistakes. These readings use various rhetorical devices to explain the concepts to their readers. In our analysis of a collection of books, online materials, and a design workshop, we identified six common explanation methods. To assess the effectiveness of these methods, we conducted two crowdsourced studies (each with $N = 125$) to evaluate their ability to teach and persuade people to make design changes. In addition to these existing methods, we brought in the idea of Explorable Explanations, which allows readers to experiment with different chart settings and observe how the changes are reflected in the visualization. While we did not find significant differences across explanation methods, the results of our experiments indicate that, following the exposure to the explanations, the participants showed improved proficiency in identifying deceptive charts and were more receptive to proposed alterations of the visualization design. We discovered that participants were willing to accept more than 60% of the proposed adjustments in the persuasiveness assessment. Nevertheless, we found no significant differences among different explanation methods in convincing participants to accept the modifications.

Index Terms—Information Visualization, Deceptive Visualization, Explorable Explanations

◆

1 INTRODUCTION

The visualization community has developed guidelines for practitioners, outlining best practices and common pitfalls to avoid. These guidelines, such as starting the Y-axis at zero when drawing a bar chart and selecting sequential color schemes for continuous variables, are frequently emphasized. In addition, numerous books written by authors such as Huff [18], Tufte [36], Cairo [3], and Monmonier [28], as well as many other articles [9, 31, 35], have addressed how poorly constructed visualizations can mislead the audience.

The major takeaway from these readings is that readers are equipped with the skills necessary to (1) identify misleading visualizations, and (2) avoid creating deceptive visualizations themselves. These two abilities encompass the dual roles of being both consumers and producers of data visualizations in everyday life. A common rhetorical strategy in these works involves presenting an example of a misleading visualization and then systematically explaining why it is deceptive. This emphasis on explanation is crucial in educating readers to recognize misleading visualizations others produced and avoid making similar mistakes in their own work.

Explaining data visualization best practices is also crucial to automatic visualization recommendation and correction systems. Early efforts in detecting visualization errors primarily focused on finding the issues within a given visualization, typically supplying an error message as the output [6, 26, 27]. While it is important for these automated systems to detect problems and potentially suggest solutions, the communication aspect is often overlooked.

To effectively integrate these diagnostic systems into existing visualization tools and assist creators, it is necessary to address the communication gap. Presenting users with clear explanations of why their visualization design may be potentially misleading and persuading them to accept the suggested fix is a challenge that must be tackled before these systems can be widely adopted and trusted by the end users.

• Leo Yu-Ho Lo, Yifan Cao, Leni Yang, and Huamin Qu are with the Hong Kong University of Science and Technology, Hong Kong. E-mail: {yhload@cse, ycaoaw@connect, lyangbb@cse, huamin@cse}.ust.hk.

In this work, we investigate practical methods for delivering explanations to readers. We begin by reviewing the literature and examining examples collected from the internet to summarize how people explain misleading visualizations to their audiences. Subsequently, we organized an explanation formulation workshop to gather ideas on elucidating deceptive visualizations and to apply the explanation techniques derived from our literature review and online sources.

Through this process, we identified six techniques for explaining misleading visualizations, which can serve as a foundation for improving the communication of data visualization best practices and enhancing the effectiveness of automatic visualization recommendation and correction systems.

In addition to the existing explanation techniques, we propose adopting Bret Victor's concept of Explorable Explanations [38] as an interactive method for conveying information to readers. Explorable Explanations enable users to engage with and investigate complex concepts, ideas, and data in an interactive and immersive manner. This approach is well-suited to address the explanatory challenges faced by readers and, potentially, visualization tool users. It also provides an engaging and interactive medium to comprehend and reassess the design choices of potentially deceptive visual representations.

We conducted two between-subject experiments on crowdsourcing platforms to evaluate the learnability and persuasiveness of five explanation techniques for addressing various visualization mistakes. In the learnability experiment, we measured participants' prior- and post-intervention performance in identifying misleading visualizations to assess the effectiveness of the methods. For the persuasiveness experiment, we evaluated the acceptance rate of the suggested visualization corrections to determine the impact of the explanation techniques on users' willingness to adopt the proposed changes.

Through this work, we aim to contribute the following to the visualization community: (1) a comprehensive compilation of existing explanation techniques for addressing misleading visualizations; (2) a prime example of Explorable Explanations focused on clarifying deceptive visualizations; and (3) two evaluation experiments assessing the learnability of identifying misleading visualizations and the persuasiveness of accepting suggested corrections.

2 RELATED WORK

Our study builds on the pioneering research on misleading visualizations, also known as deceptive visualizations. Previous studies collected the existence of misleading visualizations and coined the term "lies"

to describe them. Recent work focuses on developing algorithms to automatically detect them when the users are creating visualizations using computer programming tools. Yet, there is a significant gap in the research when it comes to assessing the effectiveness of different explanation methods in making these issues clear to users or readers.

2.1 Misleading Visualizations

The two most influential books on this topic are Huff's *How to Lie with Statistics* in 1954 [18] and Tufte's *The Visual Display of Quantitative Information* in 1983 [36]. Both books addressed the issue of misleading charts, using examples gathered from the news media of their respective times. The deceptive tactics discussed in these books remain prevalent and widely debated. Tufte's introduction of the Lie Factor, a formula for identifying misleading charts, is a unique example of a quantitative approach to the problem. Other explanation techniques featured in these classic books, such as correction and annotation, continue to be widely employed when discussing visualization pitfalls and best practices.

Monmonier authored *How to Lie with Maps* [28] from a cartographer's perspective, specifically addressing the misleading presentation of information and geographic data in maps. Jones's book, *How to Lie with Charts* [19], offers a creator's perspective on avoiding the production of deceptive charts when using spreadsheets and slideshow software. Cairo's *How Charts Lie* [3] provides a modern take on the topic, featuring recent examples of misleading techniques collected from social media and news media. While these works expand upon the issues initially discussed by Huff and Tufte by examining a broader range of charts beyond printed media, the techniques used in their explanations remain similar.

Misinformation is widespread on the internet, with data visualization often serving as a key medium for disseminating such inaccurate information. Several studies have investigated the use of visualizations in the context of misinformation online. Lo *et al.* aimed to broaden the scope of poorly constructed visualizations, addressing a wide variety of deceptive techniques that result in misleading representations and mistakes that render visualizations uninformative [23]. Lee *et al.* analyzed social media to understand the role of visualizations in social media posts and the dissemination of misinformation [20]. Lisnic *et al.* collected and examined instances of visualizations being used to spread misinformation, irrespective of the deceptiveness of the charts [22].

The deceptive impact of misleading visualizations has been consistently demonstrated across various studies. Pandey *et al.* evaluated the effects of message exaggeration/understatement and message reversal through a crowdsourced experiment. Their results revealed that participants who viewed the misleading version of the chart interpreted the data differently [31]. In another study, Correll *et al.* focused on different visual designs of truncated axes trying to mitigate the misleading effect of a truncated axis [8]. Their findings indicated that, despite attempts to mitigate the issue by hinting at the presence of a truncated axis, truncation consistently led to an exaggerated perception of the differences between bars.

Critical thinking is an important skill in data visualization literacy. Reflecting on elementary school teaching, Chevalier *et al.* stressed the necessity of embedding critical thinking within visualization literacy education [7]. Bergstrom and West, in their university course [1], gathered and deliberated on misleading instances of data visualizations, which were later compiled in their book [2]. In their hands-on teaching classroom, Lo *et al.* guided students to construct charts using visualization software and experience firsthand the inadequacy of the default axis direction in the context of ranking data, which smaller value represents better ranking. This inadvertently led to a downward trend for the ascending ranking. Subsequently, the students were tasked to rectify the chart by adjusting the chart settings [24]. Camba *et al.* investigated the effectiveness of various learning activities–such as in-class discussions, self-learning, and peer challenges–on learning to identify and accurately interpret misleading visualizations. They found that the peer challenge intervention significantly improved performance in the post-intervention test. For this activity, students were asked to experiment with the dataset and different visualization settings to "deceive" their peers and identify deceptive techniques used in their peers' visualizations. This active learning approach substantially benefited students' learning outcomes [4].

2.2 Automated Detection and Explanation of Misleading Visualizations

The prevalence and impact of misleading visualizations in modern communication are significant, especially given the widespread availability of visualization tools. Automated detection and explanations serve as important countermeasures against these issues. Such detection tools are often referred to as linters, drawing an analogy to computer program code linters that raise warnings for constructs that may be legitimate but have the potential to cause errors. This analogy is fitting for detecting visualizations that are constructed correctly but may still be misleading.

McNutt and Kindlmann proposed a linter with linting rules derived from the Algebraic Visualization Design (AVD) framework [26]. They implemented a Python-based linter for charts created using the Python charting library–matplotlib. In a course project, Zheng and Sherif developed a JavaScript-based linter for the charting library, Chart.js [40]. Vizlinter proposed by Chen *et al.*, is built with Answer Set Programming (ASP), and it checks the visualization specifications of Vega-Lite [6]. These linters are rule-based, relying on predefined rules derived from visualization best practices. Contrarily, McNutt *et al.* applied the AVD framework to perform automatic checking on visualizations without predefined rules [27].

While the primary goal of these tools is to create automated checking programs that detect poorly constructed visualizations, their textual outputs need to be more approachable for end users to understand the issues and accept the needed changes. These automatic detection tools provide only textual warnings, which may notify users but lack clear explanations of the problem and its location within the visualization.

Hopkins *et al.* proposed Visualint, a visual interface designed to highlight problematic regions of a chart and notify users that the visualization is poorly constructed [17]. Fan *et al.* suggested a pipeline that performs visualization reverse engineering [32] to extract the visualization specification from its bitmap graphics form, then checks it against the rules derived from best practices [12]. If any violations are detected, a small overlay is created on the chart for comparison with the suggested version. While these tools have limitations, such as applicability to specific chart types, reliance on expert-created rules, and potentially high computational costs, they represent pioneering efforts in exploring the feasibility of creating user interfaces to help users understand issues and correct their visualizations.

2.3 Explorable Explanations

Bret Victor proposed the concept of Explorable Explanations as a powerful rhetorical device in writing [38]. Enabled by web technology, interactive reading differs from traditional book reading by actively engaging readers with articles incorporating Explorable Explanations. Instead of using static images and precalculated numbers, writers provide small widgets that allow readers to adjust the parameters and observe how the graphs or numbers change accordingly.

Dragicevic *et al.* introduced the idea of an explorable multiverse analysis report (EMAR) for reporting experimental results [10]. By adopting Explorable Explanations, readers can explore different settings of the experiment report, increasing transparency and mitigating issues related to reporting a single analysis path.

Explorable Explanations has significant potential for explaining various concepts, regardless of their complexity. A collection of Explorable Explanations is showcased on explorabl.es [11], featuring numerous examples that cover complex mathematical and scientific concepts. One particularly inspiring example is Nicky Case's creative work on explaining the intricate concept of game theory in an interactive and engaging way with Explorable Explanations [5]. We aim to incorporate Explorable Explanations into the explanation of visualization concepts, enabling readers to learn and be persuaded to identify misleading visualizations and accept suggested design changes more effectively.

3 EXISTING FORMS OF EXPLANATIONS

In order to gain a thorough understanding of the common methods used to explain misleading visualizations, we collected examples from three distinct sources: the internet, books, and a design workshop. These examples illustrate why certain visualizations can be misleading and provide insight into the most effective ways to explain them.

3.1 Examples from the Internet

The internet offers an abundant resource for explanations concerning misleading visualizations. Numerous educational blog posts and social media discussions have been discovered in our research. We utilized the collection compiled by Lo *et al.* [23], who assembled a dataset of improperly designed visualization examples sourced from search engines and social media platforms. Within this collection, 1,143 visualization examples are tagged with at least one issue. We leverage this dataset to investigate the methods employed by individuals online when explaining misleading visualizations to others.

We first attempted to access the web pages using the URLs provided in the dataset. As of January 2023, we successfully retrieved 992 documents out of the 1,143 entries. Some documents were irretrievable due to missing web pages or removal from social media platforms. Next, we filtered out duplicate entries with identical URLs and documents lacking explanations. Following the retrieval and filtering procedures, our study incorporated 360 documents containing explanations. Notably, nearly a tenth ($9.4\%, N = 34$) of these explanations were developed as instructional materials in the form of quizzes ($4.4\%, N = 16$), slides ($3.9\%, N = 14$), or course websites ($1.7\%, N = 6$).

Two authors coded these documents using the grounded theory method (GTM) as described by Muller [29]. GTM is a research approach for exploring an unfamiliar domain—in our case, understanding how to explain misleading visualizations. The coders iteratively analyzed the examples and discussed the explanation methods' definitions. After iterations of coding and refining definitions, each tag achieved a Cohen's $\kappa > 0.7$. The coding results can be found in the supplemental materials.

The explanation methods fall into two categories: visual and textual, often complementing each other. Most explanations contain textual content ($93.9\%, N = 338$). These explanations can be as concise as directly identifying issues, such as "Look at the percentages and the length of columns!"[1] We labeled these explanation methods as **Short Text** ($48.3\%, N = 174$). Short Text explanations are predominantly found on social media platforms like Twitter, where post length is limited. Conversely, **Long Text** explanations ($45.6\%, N = 164$) are most commonly encountered in blog posts. They may involve in-depth discussions of a single issue[2] or address various visualization pitfalls using multiple examples[3].

Over half of the explanations incorporate visual aids ($55.3\%, N = 199$), with some even employing a combination of multiple techniques. The most prevalent visual aid contrasts two visualizations: a misleading version and a corrected or redrawn version. We differentiate between Correction and Redraw, as a correction only modifies chart settings, while a redraw changes the entire chart type. For instance, changing the y-axis to an appropriate range to better show the trend in data is a correction (Fig. 1a), while converting a pie chart to a bar chart is a redraw (Fig. 1b). **Corrections** ($41.7\%, N = 150$) are more common than **Redraws** ($12.8\%, N = 46$).

Another form of visual explanation involves marking the visualization to guide the audience in identifying the misleading elements in the chart. The simpler form is **Highlighting** ($5.8\%, N = 21$), and the more complex form is **Annotation** ($15\%, N = 54$). Highlighting captures the audience's attention and emphasizes a specific area on the chart by circling the crucial part that leads to misunderstanding. This can also be

---

[1] https://twitter.com/1GoldilocksZone/status/856973264526209024
[2] https://andrewpwheeler.com/2013/08/28/hanging-rootograms-and-viz-differences-in-time-series/
[3] https://flowingdata.com/2017/02/09/how-to-spot-visualization-lies/

Table 1: Classification of six explanation methods identified from online explanation examples.

Categories		Explanation Methods	Descriptions
Textual		Short Text	Provide a brief explanation of the chart's issue(s) in a few concise sentences.
		Long Text	Offer a more detailed explanation using multiple paragraphs.
Visual	Contrasting	Correction	Amend the original chart by modifying its settings.
		Redraw	Recreate the original chart using a different chart type.
	Marking	Highlighting	Highlight or emphasize a specific region in the original chart to capture the reader's attention.
		Annotation	Illustrate the issue(s) more clearly by adding annotations to the original chart.

achieved using arrows, a highlighter pen effect, or enlarging the region with a magnifying effect. Fig. 1c exemplifies circling a data point that creates a false impression of a decreasing trend. Annotations draw the audience's attention and provide additional visual aids. The annotation in Fig. 1d simplifies the comparison between two bars for the audience. There are numerous other annotation forms, such as overlaying small icons on the larger icons in pictograms to facilitate area comparisons (Fig. 1e), adding reference lines to emphasize the scale inconsistencies in the chart (Fig. 1f), or explicitly marking chart elements with in-situ text to expose the inconsistent binning sizes (Fig. 1g).

Tab. 1 presents the categorization of explanation methods and their descriptions. In addition to these six identified methods, during the coding process, we also observed instances where evidence from various sources was compiled to form arguments that debunk claims derived from the charts. This evidence could include charts from other news sources, data from different data agencies, or other evidence like a prescription drug receipt to refute an incorrect insulin price chart. We conclude that although the categorization of these six explanation methods is not exhaustive in the design space of explaining misleading visualizations, it encompasses a representative set of explanations found on the internet, providing a foundation for those seeking to develop explanations for misleading visualizations.

3.2 Explanations from Books

The first documented discussion of misleading visualizations appears in Huff's book on statistics, which laid the groundwork for subsequent conversations on the subject. Books provide a valuable corpus for analysis, so we also examined them to learn how authors explain the examples they have collected. Through related book suggestions from Amazon and Goodreads, we gathered six titles: Huff [18], Tufte [36], Monmonier [28], Jones [19], Cairo [3], and Bergstrom and West [2]. These titles contain at least one chapter related to misleading visualizations, while other titles focus on identifying misleading uses of statistics but lack dedicated chapters on visualizations [16, 21, 34, 39].

We examined the chapters related to misleading charts and categorized the explanation methods employed by the authors. We found that contrasting techniques were used most frequently. The author would present a misleading visualization example encountered in daily life, guide readers in identifying the misleading elements in the chart, and then reveal the corrected or redrawn version as a truthful data representation. Annotation is also a common technique, while Highlighting is the least common. These findings align with the results from the internet examples. We also noted Tufte's use of formulas to explain the Lie Factor, which is the only instance we found where a formula was used to explain why a chart is misleading.

3.3 Design Workshop

In addition to examples from the internet and books, we aimed to explore potential new explanation methods. We organized a design workshop with 16 participants, consisting of postgraduate students and researchers in the field of data visualization research.

Fig. 1: Examples of visual explanations from online websites and social media. (a) Explaining the poor design choice of the y-axis range by providing the corrected version of the original chart. (b) Redrawing a 3D pie chart as a bar chart or correction by removing the 3D effect. (c) Highlighting the inconsistent intervals of the last two x-axis elements in contrast to the other data points. (d) Annotating bar sizes by drawing boxes on top of the bars for easy comparison. (e) Filling the larger area with more small icons than the differences in data values suggest. (f) Drawing reference lines to emphasize the scale inconsistencies in the chart. (g) Adding in-situ text to point out the inconsistencies in bin sizes.

The workshop comprised three parts. In the first part, participants answered six questions related to various chart issues. The initial five questions each included two visualizations—one misleading and one corrected version. Participants were asked to explain why one chart was preferable to the other and to convey their explanations to a layperson from the general public. They were encouraged to use non-textual aids in their explanations. The fifth question was similar to the first four but lacked a corrected version. The final question asked participants to justify their explanations. In this part, we aimed to gather ideas about explanation methods beyond those collected from the internet and books.

The second part introduced techniques found in the internet examples to participants, aiming to provide hints or spark ideas for enhancing their initial answers. In the third part, participants revisited their explanations and attempted to improve them using techniques introduced in the second part. Lastly, they were asked to share their thoughts on the changes. This part aimed to explore the benefits of providing explanation techniques and examples for those forming explanations and rationalizing visualization design choices.

From the design workshop, we identified three new variants of explanation techniques: (1) Analogy, (2) Breakdown, and (3) Animation. Analogy is akin to Huff's chart background explanations for truncated axes [18]. Breakdown shows the decomposition of the largest group in the space next to the original chart, such as the largest group in a histogram. Animation involves rotating or morphing a chart into its correct form. As Analogy and Breakdown are drawn on the original chart to assist explanation, they may be categorized as Annotations. Animation is similar to contrasting techniques like Correction and Redraw but transforms the original chart into the corrected or redrawn version with animation. The workshop materials and coding results can be found in the supplemental materials.

Over three-quarters of the participants ($75\%, N = 12$) enhanced their explanations in part three. Participants who initially relied on textual explanations were able to enrich their explanations with various visual aids across different questions. Highlighting was the most frequently used technique ($52.5\%, N = 42$) in part one.

Through the explanations collected from the internet, books, and design workshop, we identified six major forms: (1) Short Text, (2) Long Text, (3) Correction, (4) Redraw, (5) Highlighting, and (6) Annotation. This classification informs the design of explanations that facilitate learning and persuade visualization tool users to make better

design choices. All the collected data and materials are available on OSF[4].

4 EXPLANATION DESIGNS

For the evaluation study, we selected five issues from the taxonomy proposed by Lo *et al.* [23] and five explanation methods from the formative study. Apart from the implementation for evaluation purposes, we also examined the feasibility and challenges of applying different explanation methods.

4.1 Construction of Misleading Visualizations

In selecting visualization issues to demonstrate the explanation methods, we aimed to include a diverse range of common design-related visualization pitfalls and different chart types. The design and perception stages are the most relevant from the five-stage taxonomy proposed by Lo *et al.* [23]. From these two stages, we selected three design choice issues and two perception issues with the highest occurrences in their corresponding categories. This selection covers the choice of axes, chart types, color schemes, and the design choice between 2D and 3D. We selected issues that comprise: (1) **Truncated Axis** on bar charts, where the differences between bar lengths are exaggerated due to a non-zero starting point on the y-axis. Fig. 2d provides an annotated explanation of this issue. (2) **Inappropriate Axis Range** on line charts, where the lines appear predominantly flat due to the expansive range of the y-axis (Fig. 2e). (3) **Inappropriate Use of Line Chart**, where a categorical variable is erroneously encoded on the y-axis (Fig. 2f). (4) **Ineffective Color Scheme** on choropleth maps, where a rainbow color scheme is used to represent a continuous variable (Fig. 2g). (5) **3D** pie chart, where the pie chart is presented with a perspective distortion (Fig. 2h). These issues span the most common chart types–bar charts, line charts, pie charts, and choropleths–and highlight their most prevalent issues.

In the explanation examples collected in Section 3, some explanations use charts constructed from synthetic data to better illustrate the problem. However, synthetic data may risk influencing experiment results by causing participants to question the data instead of the chart design. To avoid this, we used real data to construct the charts and clearly indicated the data sources next to the chart. We randomly picked datasets from the list of all available charts on Our World In Data [30].

Fig. 2: Examples of implemented visual explanation methods. (a) Correcting a rainbow color choropleth to a sequential color scheme. (b) Redrawing a horizontal line chart with categorical variables into a bar chart. (c) Highlighting the improperly configured axis that downplays the data variations. (d) Using lines to annotate the mismatch between bar lengths and data values in the chart. (e) Annotating to emphasize the impact of an improperly configured axis that minimizes data fluctuations. (f) Using arrows to emphasize the irrelevance of sharp inclines in a line chart with categorical variables. (g) Pointing out the confusion caused by distinct values being represented by similar hues in a rainbow color scheme. (h) Overlaying the bottom pie chart slice with the top pie chart slice that has a downplayed area caused by the 3D perspective effect.

To avoid geographical biases, country names were masked as Country A, B, C, etc., except in choropleth maps. The same practice applied to continents.

We employed Python libraries pandas [25] and Altair [37] within the Jupyter Lab environment to construct the charts, except for 3D pie charts, which we created using Microsoft Excel with VBA scripts due to a lack of available Python libraries for 3D pie charts. We constructed a total of 60 charts, which were then randomly divided into two sets. The division criteria ensured that each of the five issues was represented by three misleading charts and their respective corrected versions, thus resulting in $5 \times 3 \times 2 = 30$ charts in each set. In Experiment One, the pre-intervention test used the first set, while the post-intervention test used the second set. In Experiment Two, only the first set is used. More specifics on the experimental design are provided in Sec. 5.1 and Sec. 5.2. The datasets and scripts to construct these charts are available on OSF[4].

4.2 Explanation Methods

We implemented five designs to explain the above issues: (1) Short Text, (2) Correction or Redraw, (3) Highlighting, (4) Annotation, and (5) Explorable Explanations. **Short Text**, the simplest and most commonly used explanation method, sets the baseline for our evaluation experiment. **Correction** and **Redraw** both contrast the corrected version and the original misleading version, differing in whether the chart type is changed. Four of the five issues we studied can be corrected without changing the chart type, with the only exception of Inappropriate Use of Line Chart, which requires a redraw of the line chart to a bar chart to avoid misleading encoding of categorical variables. **Highlighting** can be done by circling chart elements related to the issues, such as the starting point of the y-axis or the legend of rainbow colors. Implementing **Annotation** requires more sophisticated thinking to convey and explain the message to readers. We learned from the collected examples from the internet and books that applied the Annotation technique, which includes using lines, arrows, short phrases, and overlapping areas to compare sizes. Examples of these designs can be seen in Fig. 2.

Explorable Explanations is a concept proposed by Bret Victor [38]. It allows readers to understand the content and interact with and challenge the writer's claim. It provides transparency, enables exploration, builds confidence and trust in the arguments, and encourages readers to form their own ideas and actively test them against the writer's message. Throughout the readings, there are controllable widgets that allow readers to tweak and see the effect on the numbers or graphics. We hypothesize that Explorable Explanations is a good match for explaining misleading visualizations. They enable readers not only to accept visualization best practices but also to explore different parameter settings

and see how they affect the charts.

We implemented explorable explanations through sliders and radio buttons for the five issues, enabling readers to interact with different parameter settings and see their effects on the charts. For Truncated Axis and Inappropriate Axis Range, readers can move sliders to change the axis start or end points. The same sliders apply to 3D pie charts, allowing readers to change the perspective angle. For choropleths, readers can change the color scheme between sequential, diverging, and rainbow colors. For Inappropriate Use of Line Chart, readers can configure the line chart with different item ordering, orientation, and chart types between line chart and bar chart. Fig. 3 shows examples of implemented Explorable Explanations.

5 EVALUATION EXPERIMENTS

We aim to evaluate the effectiveness of different explanation methods in two aspects: (1) learning to identify visualizations that violate visualization guidelines, and (2) persuading visualization tool users to choose a design that follows the visualization guidelines. We designed two experiments and conducted them on prolific.ac, a crowdsourcing platform focused on studies rather than the more general crowdsourcing platform provided by Amazon. The recommended hourly rate on the Prolific platform is GBP 9 (~USD 11).

For both studies, we require participants to be fluent in English, have no colorblindness, and maintain an approval rate of 98% or higher. On the Prolific platform, 75,552 workers meet these criteria, accounting for 62.2% of the total 121,407 available workers. The criterion of being free from colorblindness ($70.4\%, N = 85,445$) is the major factor reducing the suitable worker pool. Despite that, this criterion is necessary to exclude variations that may affect the results related to color scheme design choices. Besides these criteria, we do not impose further restrictions on gender, age, geographic location, or education level.

Both studies begin with a consent form, a background information form, and a subjective literacy assessment form [14]. The subjective graph literacy (SGL) assessment consists of ten questions, each asking participants to rate their level of competence in chart reading and reliance on graphical information on a six-point scale. The assessment is empirically tested to have a strong positive correlation with the objective graph literacy (OGL) [13] score.

5.1 Experiment One: Learning Effects of Explanations

We are interested in testing the effectiveness of explanations in educating readers to identify charts that violate visualization guidelines. In this experiment, we have set up five conditions: (1) Short Text, (2) Short Text + Highlighting, (3) Short Text + Annotation, (4) Short Text +

a. Slider to adjust the starting point of the y-axis

b. Two-sided slider to adjust the starting point and ending point of the y-axis

c. Radio buttons to modify the chart settings

d. Radio buttons to change color schemes

e. Slider to change the tilt angle of the pie chart

Fig. 3: Examples of implemented Explorable Explanations. (a) By adjusting the slider, readers can experiment with different values of the y-axis starting point to see the changes in the bar lengths. (b) Readers can set the y-axis range by adjusting the two-sided slider to observe the changes in the line chart. (c) The radio buttons on the right let readers experiment with different chart settings. (d) Readers can choose different color schemes by clicking the radio buttons. (e) Readers can change the tilt angle of the pie chart to see the changes in the slice area.

Correction, and (5) Short Text + Correction + Explorable Explanations. The Short Text condition serves as a baseline for comparison with the Highlighting, Annotation, and Correction conditions. Condition 5 requires combining Explorable Explanations with Correction because the corrected version of the chart is used to guide participants to explore the explanation.

5.1.1 Experiment Procedures

The study is conducted using a between-subject design in three phases, including a pre-intervention testing phase, an intervention phase, and a post-intervention testing phase.

After collecting the basic information and completing the SGL assessment, participants will first perform a test to gauge their ability to distinguish between misleading visualizations and those that accurately represent data. The test comprises 30 visualizations, including three misleading charts for each of the five issues and their respective corrected versions, *i.e.*, the first set described in Sec. 4.1. The order of the charts is randomized across participants. Participants answer the question with one of the options "misleading," "accurate," or "I am not sure." The screenshots of a sample test are included in the supplemental materials and are available on OSF[4].

In the intervention phase, participants view a series of five explanations constructed according to their assigned conditions. Participants are asked to rate the helpfulness of the explanation and elaborate on their rating. Fig. 4 shows the screenshot of the participants' interface in Condition 5. The order of the explanations on different issues is randomized across participants. After reading the explanations, participants are required to complete the post-intervention test, which consists of a different set of 30 visualizations with a similar composition to the pre-intervention test, *i.e.*, the second set described in Sec. 4.1. Upon completion, participants are thanked and given an opportunity to report any comments or issues they encountered during the experiment.

5.1.2 Hypotheses

In Experiment One, we want to test the effectiveness of explanations on participants' ability to identify charts that violate visualization guidelines. We hypothesize that the intervention, *i.e.*, exposure to the explanations, has an effect on participants' ability to distinguish misleading charts. Therefore, our first hypothesis is: **Hypothesis 1.1 The correctness in identifying misleading charts and accurate charts is higher in the post-intervention phase, regardless of the explanation method.**

Secondly, Explorable Explanations has demonstrated its effectiveness in illustrating the linkage between variable parameters and their impact on outcomes. In the context of visualizations, these parameters involve design decisions such as the start and end points of the axis, or the color scheme applied in the chart. We posit that Explorable Explanations can also be used effectively to explain visualization guidelines. Hence, our second hypothesis is: **Hypothesis 1.2 The Explorable Explanations condition has a greater effect on improving the correctness in the post-intervention phase.**

5.1.3 Participants

For each condition, we recruited 25 participants on the Prolific platform according to the criteria described at the beginning of Sec. 5. In total, 125 participants took part in the experiment. Among the participants, 57.6% identified as male, 40.8% as female, and 1.6% as other. The reported age ranged from 19 to 54 ($age_{median} : 24, age_{mean} : 26.2, age_{std} : 6.06$). The reported highest completed education levels were Bachelor's degree (41.6%), some college credit (23.2%), high school graduation (21%), Master's degree (12.8%), and Doctorate's degree (1.6%). The most common occupation was student (40%, $N = 50$); the rest were mostly white-collar workers, except a cooking assistant, receptionist, warehouse worker, dentist, and military personnel. The participants' nationalities were mainly from Europe (67.2%), with the remainder from Africa (18.4%), North America (9.6%), South America (2.4%), and Asia (2.4%). The SGL assessment (10-item on a six-level scale) had a median of 4.4 ($SGL_{mean} : 4.32, SGL_{std} : 0.86, SGL_{min} : 2.5, SGL_{max} : 5.9$). The median completion time for Experiment One was 23.78 minutes. The study was advertised to workers as a 25-minute study with a compensation of GBP 3.75 (~USD 4.63), translating to an effective hourly rate of GBP 9 (~USD 11). Two participants suspected of answering randomly (15 ± 1 out of 30 correct in both prior- and post-intervention tests) were checked, and their submissions were removed from the study.

5.1.4 Results

The median number of correctly answered questions on the pre-intervention test was 19 out of 30 (63.3%, $correct_{mean} : 18.69, correct_{std} : 3.38, correct_{min} : 11, correct_{max} : 27, N = 125$), while the post-intervention median was 27 (90%, $correct_{mean} : 26.02, correct_{std} : 3.44, correct_{min} : 16, correct_{max} : 30, N = 125$). An ANOVA test comparing the two measures rejected the null hypothesis for hypothesis 1.1, indicating a significant positive change in partici-

Fig. 4: Screenshot of Experiment One Condition 5 (Short Text + Correction + Explorable Explanation). The other conditions are similar to this interface. Condition 1 has only the text explanation. Conditions 2 and 3 have the text explanations and Highlighting or Annotation but not the correction. Condition 4 has the same interface but not the explorable view.

Fig. 5: Confidence intervals of correct answers in the prior- and post-intervention tests. Participants performed significantly better after reading any of the explanations. However, we did not observe significant differences across conditions.

pants' ability to correctly identify misleading and accurate charts in the post-intervention test ($F(1, 248) = 288, p < 0.0001$). Fig. 5 shows the confidence interval plots for each condition and the overall results.

Additionally, we observed the median precision, recall, and F1 scores for the pre-intervention phase as 0.71, 0.53, and 0.61, respectively, and the post-intervention phase as 0.92, 0.93, and 0.91. In the pre-intervention test, despite instructions specifying a 15 : 15 ratio between misleading and accurate visualizations, participants were more likely to identify charts as accurate rather than misleading, with a ratio of 10.8 : 17.4, suggesting they only labeled charts as misleading when confident in their decision. This ratio shifted to 15.1 : 14.4 in the post-intervention test.

We conducted post hoc pairwise t-tests with Bonferroni correction between Condition 5 (Text + Correction + Explorable Explanations) and other conditions. There was a significant difference when compared to Condition 3 (Text + Annotation, $p : 0.008$). However, no significant differences were observed when compared with other conditions. Therefore, the null hypothesis of hypothesis 1.2 is not rejected.

Aside from Condition 5, it is worth noting that Condition 3 (Text + Annotation) showed significantly lower improvement compared to other conditions with visual explanations: vs. Condition 2 (Text + Highlighting, $p : 0.066$), vs. Condition 4 (Text + Correction, $p : 0.020$), and vs. Condition 5 (Text + Correction + Explorable Explanations, $p : 0.008$). Despite that, no significant difference was found when compared to baseline Condition 1 (Text Only, $p : 1.0$).

5.2 Experiment Two: Persuasive Effects of Explanations

One of the goals of this study is to inform the future development of visualization tools by suggesting chart fixes that align with visualization guidelines. We designed Experiment Two to simulate a situation where participants are persuaded to abandon their choice of a misleading chart and switch to one that aligns with visualization guidelines. In this experiment, we established five conditions: (1) Correction only, (2) Correction + Short Text, (3) Correction + Short Text + Highlighting, (4) Correction + Short Text + Annotation, and (5) Correction + Short Text + Explorable Explanations. Since we suggest a fix to users, the base condition shows only the corrected version. Condition 2 adds on the base condition with Short Text explanations. Conditions 3, 4, and 5 provide further visual explanations.

5.2.1 Experiment Procedures

Experiment Two employs a between-subject design with only two phases. Same as Experiment One, participants begin by completing basic information and an SGL assessment. The first phase includes 15 this-or-that questions, asking participants to choose one of two presented chart designs as the better one. One chart in each pair exhibits one of the five misleading issues, while the other is a corrected version. Each issue has three pairs of charts in the question set. The visualization pairs and their order within the pair are randomized.

Based on answers from the first phase, the chosen misleading charts are presented again in the second phase, accompanied by explanations of why the selected charts are misleading and why the alternative charts are better. Participants are asked to accept or reject the suggestion,

Fig. 6: Screenshot of Experiment Two Condition 5 (Correction + Short Text + Explorable Explanations). The other conditions are similar to this interface. Condition 1 has only the correction. Condition 2 has the correction and the text explanations. Conditions 3 and 4 have the same interface, replacing the explorable view with Highlighting or Annotation.

or choose neither options. They must also explain their choice and rate their confidence on a five-option Likert scale. Fig. 6 shows the screenshot of the interface displayed to the participants in Condition 5. After completing the experiment, participants are thanked and allowed to report any comments or issues they encountered.

5.2.2 Hypotheses

Experiment Two aims to test whether participants might change their minds and accept the suggested chart version. We hypothesize that providing an explanation alongside the corrected chart is more convincing than presenting only the corrected chart. Thus, our first hypothesis is: **Hypothesis 2.1 The participants in conditions with more than just the corrected chart will have a higher acceptance rate for the suggested chart.**

Secondly, similar to Experiment One, we believe that people's strong preferences for rainbow colors and 3D pie charts make it difficult to persuade them to choose less visually appealing monotonic sequential colors and 2D pie charts. Our second hypothesis is: **Hypothesis 2.2 The acceptance rates across different issue types will exhibit significant differences regardless of explanation methods.**

Lastly, we have observed the effectiveness of Explorable Explanations in clarifying complex concepts across various subjects, from mathematics to social sciences. We contend that applying Explorable Explanations will help communicate visualization guidelines to participants and persuade them to accept the suggested chart design. Our third hypothesis is: **Hypothesis 2.3 The acceptance rate in the Explorable Explanations condition will be higher than in other conditions across different issue types.**

5.2.3 Participants

For each of the five conditions, we recruited 25 participants through the Prolific platform based on the criteria described at the beginning of Sec. 5. Participants who participated in Experiment One were excluded from the worker pool of Experiment Two. In total, 125 participants took part in the experiment.

Participants in the second experiment were 48.8% male, 50.4% female, and 0.8% other, and their age ranged from 18 to 55 (age_{median} : 25, age_{mean} : 27.2, age_{std} : 7.78). The reported highest completed education levels included Bachelor's degree (47.2%), some college credit (24%), high school graduation (14.4%), Master's degree (11.2%), and Doctorate's degree (3.2%). The most common occupation was student (31%, N=38), with the remainder mostly comprising white-collar workers, with some exceptions such as homemakers and bus drivers. Participants' nationalities were mainly European (50.4%) and African (28.8%), with the remainder being North American (16.8%), Asian (2.4%), South American (0.8%), and Oceanic (0.8%). The SGL assessment (10-item on a six-level scale) showed a median of 4.5 (SGL_{mean} : 4.4, SGL_{std} : 0.81, SGL_{min} : 2.0, SGL_{max} : 6.0). The median completion time for Experiment Two was 13.57 minutes. The study was advertised as a 15-minute study with a GBP 2.25 (~USD 2.78) compensation, yielding an effective hourly rate of GBP 9 (~USD 11).

5.2.4 Results

The median number of correctly answered questions in the first phase was 11 out of 15 (73.3%, $correct_{mean}$: 10.66, $correct_{std}$: 2.68, $correct_{min}$: 4, $correct_{max}$: 15, $N = 125$). This result is higher than the median in Experiment One, which was 63.3%. When presented with two versions of the charts, participants were able to distinguish the misleading chart better and choose the correct one, as shown by the ANOVA test ($F(1,248) = 21.67, p < 0.001$).

Since the number of incorrect answers in phase one varied among participants, the number of visualization pairs presented to each participant also differed. The same applied to issue types. While 542 explanations were presented across all participants, only 17 were related to Inappropriate Axis Range, and 26 were concerned with Inappropriate Use of Line Chart. Due to the limited sample size, we excluded these two issues from the following analysis.

We calculated the acceptance rate as $\frac{A}{S}$, where A is the number of accepted suggestions, and S is the number of suggestions presented to

Fig. 7: Confidence intervals of the acceptance rate in phase 2 of Experiment Two, *i.e.*, the percentage of accepting the suggested correct version of the chart after viewing the explanation.

the participant. Participants who received no suggestions, *i.e.*, those who answered correctly on all nine questions, were not counted. In total, 17 participants scored full marks: 4 for Condition 1, 3 for Conditions 2 and 5, 2 for Condition 3, and 5 for Condition 4. The overall acceptance rate across participants was 60.4% ($N = 108$).

A post hoc pairwise t-test with Bonferroni correction revealed no significant differences across conditions. The confidence interval plot in Fig. 7 showed a high data variation. It remains to be seen whether adding explanations on top of presenting only the corrected version is more convincing or has no effect. Therefore, we cannot reject the null hypothesis of Hypothesis 2.1. However, the Annotation condition exhibited a high mean acceptance rate, which could be a hint for further investigation.

Across different issues, Truncated Axis had a significantly higher acceptance rate than the other two issues, vs. Inappropriate Color Scheme ($p : 0.007$) and vs. 3D pie charts ($p : 0.046$). Since we analyzed only three issues after excluding the two issues of limited sample size, we cannot determine whether participants were more convinced by the explanations on Truncated Axis or were simply reluctant to accept the corrected versions. Nonetheless, there was a difference in the acceptance rate across issue types, leading us to reject the null hypothesis of Hypothesis 2.2.

Since Hypothesis 2.1 is not supported, and Hypothesis 2.3 is a stronger claim than Hypothesis 2.1, Hypothesis 2.3 is also not supported. We cannot conclude that Explorable Explanations has a stronger persuasive power of convincing users to change their design choices.

The experiment data and processing scripts can be found on OSF[4].

6 DISCUSSION AND FUTURE WORK

Our data found little to no correlation between the SGL assessment score and the accuracy of identifying misleading charts, neither in overall accuracy nor accuracies of each individual issue type ($0 < r < 0.2$). The questions in the SGL assessment focus on the ability to interpret accurate charts but do not assess the participants' ability to question the validity of the charts. In the OGL assessment [13], there are questions where the correct answer is "cannot infer from the chart," but this aspect of critical chart reading is neglected in SGL. Recently, Ge *et al.* have compiled a set of 45 assessment questions designed to evaluate critical thinking abilities when confronted with misleading visualizations [15]. We recommend that future studies on this topic consider modifying the SGL assessment by including one or two questions to address the weakness of SGL.

In the second phase of Experiment Two, we collected participants' justifications for why they chose to keep their selection of a misleading chart over the accurate one. We coded the responses and analyzed the results. 58 out of 186 responses rejected the correction on Truncated Axis. Participants found that the lengths of the bars without truncating the axis were too similar and difficult to compare ($79.3\%, N = 46$). Even though they were aware of the problem of exaggeration, they still preferred the chart with a truncated axis. Ritchie *et al.* also investigated the value of the truncated axis over the one that shows the whole axis. They designed an interaction technique to allow readers to switch between the whole and the truncated axis [33].

61 out of 138 responses rejected the correction on Inappropriate Color Scheme. Participants reported that sequential colors were difficult to read, and rainbow colors offered better distinguishability and chart readability ($90.2\%, N = 55$). One participant rejected the correction: "It is much easier to differentiate between rainbow colours than different

shades of one colour." There were also responses stating that rainbow colors were helpful to colorblind people ($6.6\%, N = 4$), even though rainbow colors are, in fact, very difficult to read for people with color weaknesses. 63 out of 175 responses rejected the correction on 3D pie charts. Participants preferred 3D pie charts for aesthetic reasons ($41.3\%, N = 26$) and believed that since the values were labeled, readers could always refer to the labels and not be misled by them ($34.9\%, N = 22$). One participant wrote, "Your explanation might be true. However, the numbers are clearly marked."

Some justifications are valid and worth considering when constructing explanations to address misunderstandings. The preference for high contrast (between colors or bar length) over accurate interpretation may be a design trade-off from the participants' perspective. Providing better alternatives and explaining the implications to users or readers can help them make more informed design decisions. Although the guidelines for visualization design are generally considered best practices, their application is subject to the context and the specific purpose of the visualization. If deviations from these guidelines can be justified logically, they can benefit the task in question. The active field of research aiming to understand human perception and the interpretation of visualizations is persistently questioning current practices, unveiling their shortcomings, and improving both visualization tools and guidelines.

The main idea of this work is inspired by the Explorable Explanations developed for different topics across various fields and conjecturing that applying this approach to explain visualization guidelines could be valuable. The work by Hopkins *et al.* implemented and evaluated the Highlighting explanation method [17]. The experiment result is similar to our results in this study, showing no significant improvement when adding visual explanations on top of text explanations. The work by Fan *et al.* implemented the Correction explanation, and their evaluation supported the notion that people with access to explanations perform better at identifying misleading visualizations [12]. Though, they did not compare it with the base case of text-only explanations. It remains an open question: Is text explanation alone enough to explain misleading visualizations? We must admit that we did not find evidence to reject this idea. Imagine that if the classic books by Huff [18] and Tufte [36] had used only text explanations, they might not have inspired many other authors to publish titles or blog posts on the topic. Consequently, the public might have become even less aware of the pitfalls of misleading charts.

One of the primary objectives of this study is to inform the future development of visualization recommendation systems and linting systems regarding the design of user interfaces when the system attempts to suggest a better alternative visualization to the user. Although our experiment results did not conclude that Explorable Explanations has greater persuasive power over other explanation methods, it is worth noting that the design space of Explorable Explanations is much larger and highly dependent on the explanation authors' creativity. Further studies may yield more conclusive results.

7 CONCLUSION

This study identified six distinct explanation techniques through an extensive review of online resources, books, and a design workshop. We applied these methods to five chart-related issues, incorporating the concept of Explorable Explanations. Our two crowdsourced evaluation experiments revealed that exposure to explanations enhanced participants' ability to recognize misleading charts. However, no significant differences were observed among the various explanation techniques. We discovered that participants were inclined to accept more than 60% of the proposed adjustments in the persuasiveness assessment.

Nevertheless, we found no significant differences among the explanation methods in convincing participants to accept the modifications. It remains a challenge to develop more effective explanation techniques in educating and convincing the general audience on the application of visualization best practices. At a minimum, they need to be made aware of the potential pitfalls and benefits when they decide not to follow these guidelines.

REFERENCES

[1] C. T. Bergstrom and J. West. Calling bullshit in the age of big data. *Calling Bullshits. Data Reasoning in a Digital World. Available online at: http://callingbullshit.org*, 2017. 2

[2] C. T. Bergstrom and J. D. West. *Calling bullshit: The art of skepticism in a data-driven world*. Random House Trade Paperbacks, 2021. 2, 3

[3] A. Cairo. *How Charts Lie: Getting Smarter about Visual Information*. WW Norton & Company, 2019. 1, 2, 3

[4] J. D. Camba, P. Company, and V. L. Byrd. Identifying deception as a critical component of visualization literacy. *IEEE Computer Graphics and Applications*, 42(1):116–122, 2022. 2

[5] N. Case. The evolution of trust. http://ncase.me/trust, 2017. (Accessed on 03/31/2023). 2

[6] Q. Chen, F. Sun, X. Xu, Z. Chen, J. Wang, and N. Cao. Vizlinter: A linter and fixer framework for data visualization. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):206–216, 2021. 1, 2

[7] F. Chevalier, N. H. Riche, B. Alper, C. Plaisant, J. Boy, and N. Elmqvist. Observations and reflections on visualization literacy in elementary school. *IEEE Computer Graphics and Applications*, 38(3):21–29, 2018. 2

[8] M. Correll, E. Bertini, and S. Franconeri. Truncating the y-axis: Threat or menace? In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pp. 1–12, 2020. 2

[9] M. Correll and J. Heer. Black hat visualization. In *Workshop on Dealing with Cognitive Biases in Visualisations (DECISIVe), IEEE VIS*, 2017. 1

[10] P. Dragicevic, Y. Jansen, A. Sarma, M. Kay, and F. Chevalier. Increasing the transparency of research papers with explorable multiverse analyses. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pp. 1–15, 2019. 2

[11] Explorable explanations. https://explorabl.es. (Accessed on 03/31/2023). 2

[12] A. Fan, Y. Ma, M. Mancenido, and R. Maciejewski. Annotating line charts for addressing deception. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pp. 1–12, 2022. 2, 9

[13] M. Galesic and R. Garcia-Retamero. Graph literacy: A cross-cultural comparison. *Medical Decision Making*, 31(3):444–457, 2011. 5, 9

[14] R. Garcia-Retamero, E. T. Cokely, S. Ghazal, and A. Joeris. Measuring graph literacy without a test: A brief subjective assessment. *Medical Decision Making*, 36(7):854–867, 2016. 5

[15] L. W. Ge, Y. Cui, and M. Kay. Calvi: Critical thinking assessment for literacy in visualizations. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pp. 1–18, 2023. 9

[16] T. Harford. *The data detective: Ten easy rules to make sense of statistics*. Penguin, 2022. 3

[17] A. K. Hopkins, M. Correll, and A. Satyanarayan. Visualint: Sketchy in situ annotations of chart construction errors. In *Computer Graphics Forum*, vol. 39, pp. 219–228. Wiley Online Library, 2020. 2, 9

[18] D. Huff. *How to lie with statistics*. WW Norton & Company, 1954. 1, 2, 3, 4, 9

[19] G. E. Jones. *How to lie with charts*. Sybex San Francisco, 1995. 2, 3

[20] C. Lee, T. Yang, G. D. Inchoco, G. M. Jones, and A. Satyanarayan. Viral visualizations: How coronavirus skeptics use orthodox data practices to promote unorthodox science online. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pp. 1–18, 2021. 2

[21] D. J. Levitin. *A field guide to lies: Critical thinking in the information age*. Penguin, 2016. 3

[22] M. Lisnic, C. Polychronis, A. Lex, and M. Kogan. Misleading beyond visual tricks: How people actually lie with charts. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pp. 1–21, 2023. 2

[23] L. Y.-H. Lo, A. Gupta, K. Shigyo, A. Wu, E. Bertini, and H. Qu. Misinformed by visualization: What do we learn from misinformative visualizations? In *Computer Graphics Forum*, vol. 41, pp. 515–525. Wiley Online Library, 2022. 2, 3, 4

[24] L. Y.-H. Lo, Y. Ming, and H. Qu. Learning vis tools: Teaching data visualization tutorials. In *2019 IEEE Visualization Conference (VIS)*, pp. 11–15. IEEE, 2019. 2

[25] W. McKinney et al. pandas: a foundational python library for data analysis and statistics. *Python for High Performance and Scientific Computing*, 14(9):1–9, 2011. 5

[26] A. McNutt and G. Kindlmann. Linting for visualization: Towards a practical automated visualization guidance system. In *VisGuides: 2nd Workshop on the Creation, Curation, Critique and Conditioning of Principles and Guidelines in Visualization*, 2018. 1, 2

[27] A. McNutt, G. Kindlmann, and M. Correll. Surfacing visualization mirages. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pp. 1–16, 2020. 1, 2

[28] M. Monmonier. *How to lie with maps*. University of Chicago Press, 1991. 1, 2, 3

[29] M. Muller. Curiosity, creativity, and surprise as analytic tools: Grounded theory method. In *Ways of Knowing in HCI*, pp. 25–48. Springer, 2014. 3

[30] Charts - our world in data. https://ourworldindata.org/charts. (Accessed on 03/31/2023). 4

[31] A. V. Pandey, K. Rall, M. L. Satterthwaite, O. Nov, and E. Bertini. How deceptive are deceptive visualizations? an empirical analysis of common distortion techniques. In *Proceedings of the 2015 CHI Conference on Human Factors in Computing Systems*, pp. 1469–1478, 2015. 1, 2

[32] J. Poco and J. Heer. Reverse-engineering visualizations: Recovering visual encodings from chart images. In *Computer graphics forum*, vol. 36, pp. 353–363. Wiley Online Library, 2017. 2

[33] J. Ritchie, D. Wigdor, and F. Chevalier. A lie reveals the truth: Quasimodes for task-aligned data presentation. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pp. 1–13, 2019. 9

[34] S. Stephens-Davidowitz. *Everybody lies: What the internet can tell us about who we really are*. Bloomsbury Publishing, 2018. 3

[35] D. A. Szafir. The good, the bad, and the biased: Five ways visualizations can mislead (and how to fix them). *Interactions*, 25(4):26–33, 2018. 1

[36] E. R. Tufte and P. R. Graves-Morris. *The visual display of quantitative information*. Graphics press Cheshire, CT, 1983. 1, 2, 3, 9

[37] J. VanderPlas, B. Granger, J. Heer, D. Moritz, K. Wongsuphasawat, A. Satyanarayan, E. Lees, I. Timofeev, B. Welsh, and S. Sievert. Altair: interactive statistical visualizations for python. *Journal of Open Source Software*, 3(32):1057, 2018. 5

[38] B. Victor. Explorable explanations. http://worrydream.com/ExplorableExplanations, 2011. (Accessed on 03/31/2023). 1, 2, 5

[39] C. Wheelan. *Naked statistics: Stripping the dread from the data*. WW Norton & Company, 2013. 3

[40] W. Zheng and Y. Sherif. Visualization linter. https://www.cs.ubc.ca/~tmm/courses/547-19/projects/wei-youssef/finalreport.pdf Course Project, 2019. (Accessed on 03/31/2023). 2