

Du Bois Wrapped Bar Chart: Visualizing Categorical Data with Disproportionate Values

Alireza Karduni¹, Ryan Wesslen¹, Isaac Cho², Wenwen Dou¹

¹University of North Carolina at Charlotte, <akarduni, rwesslen, wdou1>@uncc.edu

²North Carolina A&T State University, icho@ncat.edu

ABSTRACT

We propose a visualization technique, Du Bois wrapped bar chart, inspired by work of W.E.B Du Bois. Du Bois wrapped bar charts enable better large-to-small bar comparison by wrapping large bars over a certain threshold. We first present two crowdsourcing experiments comparing wrapped and standard bar charts to evaluate (1) the benefit of wrapped bars in helping participants identify and compare values; (2) the characteristics of data most suitable for wrapped bars. In the first study (n=98) using real-world datasets, we find that wrapped bar charts lead to higher accuracy in identifying and estimating ratios between bars. In a follow-up study (n=190) with 13 simulated datasets, we find participants were consistently more accurate with wrapped bar charts when certain category values are disproportionate as measured by entropy and H-spread. Finally, in an in-lab study, we investigate participants' experience and strategies, leading to guidelines for when and how to use wrapped bar charts.

Author Keywords

bar chart; graphical perception; user study; evaluation; Mechanical Turk; crowdsourcing; Information visualization

CCS Concepts

•Human-centered computing → Empirical studies in visualization; Information visualization;

INTRODUCTION

The concept of a bar chart was first introduced by French scientist Nicole Oresme in the 14th century [4, 15]. However, many sources attribute the wide adoption of bar charts to the 1786 seminal work by William Playfair on “Exports and Imports of Scotland to and from different parts for one Year from Christmas 1780 to Christmas 1781” [26, 31].

The modern definition of a bar chart is “a chart that presents categorical data with rectangular bars with heights or lengths proportional to the values that they represent” [29]. Bar charts are now considered one of the most popular and prolific visualization techniques for communicating categorical values [27, 35]. However, certain data characteristics such as disproportionately large and small values make performing certain

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI '20, April 25–30, 2020, Honolulu, HI, USA.

© 2020 Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 978-1-4503-6708-0/20/04 ...\$15.00. <http://dx.doi.org/10.1145/3313831.3376365>

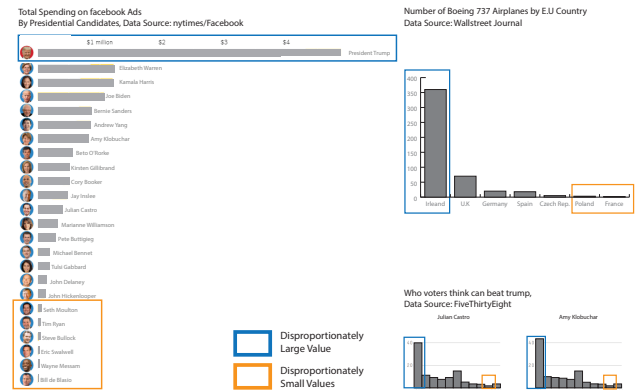


Figure 1. Three recreated examples of bar charts based on data and charts in news articles [21, 7, 9]. All three charts are characterized by one extremely disproportionate value, leaving the smaller values difficult to estimate and compare.

tasks with bar charts challenging. Figure 1 shows three examples we identified from news articles [21, 7, 9] that use bar charts to visualize categorical values¹. All three examples are characterized by one disproportionately large value, which induces a significant white space-to-data ratio in the chart’s plane. Estimating the smallest values and the ratio of largest to smallest values is challenging with charts like these. To address these challenges, we propose a new technique called a “wrapped bar chart”.

Our work is inspired by the work of William Edward Burghardt “W. E. B.” Du Bois, a sociologist, historian, activist, and author who was also a prolific graphic designer [1]. His work for the 1900 Exposition Universelle in Paris utilized hand-drawn data visualizations to focus on the African American population in the 1890’s South [3]. Du Bois explored many different types of visualization techniques including using bar charts to highlight inequities for African Americans in Georgia. For example, African Americans worked disproportionately in agriculture occupations such as laborers, farmers, and planters but rarely in manufacturing and professional occupations like engineers, masons, merchants, and barbers. This extreme difference made drawing African-American occupation counts with a standard bar chart to be ineffective, especially when comparing the largest occupation (bar) to the smallest occupation (bar). To address this problem, he introduced an innovative solution to his bar charts: to wrap the tallest bars around themselves, allowing more room to

¹We recreated the charts from the original data due to copyright issues.

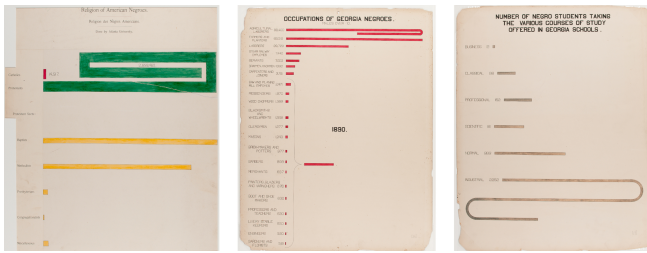


Figure 2. Three examples of wrapped bar charts designed by Du Bois [1]. Images assessed from Library of Congress. Images courtesy of Library of Congress, Prints & Photographs Division, LC-USZ62-1234

show smaller bars. Figure 2 presents three charts by Du Bois that include a “wrapped” bar for its largest data value.

Present-day visualization practitioners attempt to deal with the problem of disproportionate categorical data in bar charts through ad hoc solutions that have known issues. For example, one common solution is “breaking the axis” by interrupting the vertical axis of a chart at a specific point through a break. Similarly, cut-off bars limit the upper bound of the axis. However, these techniques are known to mislead users [8]. While the two aforementioned techniques use a linear scale, other studies have analyzed scientific data with a non-linear scale (e.g. logarithmic) when data values cover a large range of magnitudes [6]. However, to our knowledge Du Bois’ wrapping technique has not been studied by the visualization community to evaluate its effectiveness. Our work is to explore the potential and usages of Du Bois’ intriguing innovation. To this aim, we first introduce a technique for developing wrapped bar charts in web-based visualizations. Second, we conduct three experiments to explore the limitations, potentials and benefits of wrapped bar charts in helping users conduct data analysis tasks in datasets with large variance between data values. Our paper makes the following contributions:

- We design and develop the Du Bois wrapped bar chart² inspired by Du Bois’ work for web-based visualizations using D3.js.
- We conduct two online crowd-sourced experiments for comparing participants’ performance with wrapped bar charts versus standard bar charts in identification and ratio estimation tasks, using accuracy and time spent to complete each task as metrics of performance.
- We conduct an in-lab focus group study that focuses on gathering information about participants’ strategies when using wrapped bar chart, as well as collecting suggestions to improve the wrapped bar design.

In the following sections, we review related work, design of wrapped bar charts, experiment design, discussion of results, and conclude with limitations and future works.

²The similar term “wrapped bar graph” has been used for a different visualization technique. Stephen Few [17] introduced this term for a chart that splits the sorted bars in a horizontal bar chart into multiple columns to eliminate the need for scrolling [24]. In contrast, the Du Bois wrapped bar chart wraps disproportionately large bars so that small values are discernible.

RELATED WORK

Graphical Perception Studies of Bar Charts

Many empirical studies have been conducted on understanding users’ perception of the visual encoding of bar charts [13, 19]. Findings from these experiments provide design guidelines and considerations for applying bar charts to various tasks such as difference comparison and proportion estimation. The bar chart experiment by Cleveland & McGill studied participants’ estimate of the proportion of the lengths of two bars [13]. The experiment included 10 different pairs of bars with proportion values ranging from 17.8% to 82.5%. Participants’ responses were measured by the average log absolute error - the difference between a participant’s answer and the true percentage value. The analysis results revealed that the accuracy was affected by multiple factors: the responses were more inaccurate with the stacked bar charts than standard bar charts; the responses were most inaccurate for true percents around 60% - 80%; and the accuracy decreased with increasing space between the two bars.

As an effort to study crowdsourcing for perceptual experiments, Heer & Bostock conducted an approximate replication of Cleveland & McGill’s bar chart study with slight modifications on the true percents [19]. The new results are consistent with the findings from the original study although the log absolute error results are better in the new study, possibly due to the true percents being rounded to whole numbers instead of numbers with decimals. More recently, through a series of follow-up experiments, Talbot et al. [30] aimed to explore the results from Cleveland & McGill’s experiment in order to explain the sources of bar chart interpretation error. Their findings – including evidence that shorter bars are more difficult to compare – suggest that additional studies are needed to provide a more complete understanding of what impacts (bar) chart perception. Talbot et al. also found that the separation (gap) between bars increase bar comparison difficulty but the effect of intervening distractor bars are small. More recently, Zhao et al. [35] conducted experiments to evaluate how the perception of a bar changes based on the heights of its neighboring bars and found the neighborhood effect does exist.

Our study is inspired by the design and findings from the aforementioned graphical perception experiments. More specifically, the tasks in our study include the ratio estimate and identification tasks that require comparison of bars appearing both similar and orders-and-magnitude different in heights. A notable difference between the ratio estimation task in this study and previous studies on bar charts is that we chose to ask participants to calculate the ratio of largest/smallest bars as opposed to the percentage of smallest/largest bars. While the smallest percentage (small/large) evaluated by Cleveland et al. and Heer et al. was 17.8%, our study would involve estimating percentages as small as 0.2% and 0.24% if framed as small/large estimation seen in the previous studies.

Evaluating Alternative Bar Chart Designs

Multiple studies have evaluated alternative bar chart designs. Skau et al. [25] introduced different illustrative embellishments which are design alternatives of a (single-series) bar

chart. Skau et al.'s alternative bar chart designs drew inspiration from visual embellishments [5] which provide non-linguistic rhetorical figures that are used frequently in the visual/performing arts, advertisements, graphical user interfaces, etc. Srinivasan et al. [28] introduced design variants of the multi-series bar chart. Both studies evaluated the performance on the task of comparing values between bars. Both studies added additional information to single- and multi-series bars such as embellishments [25] and difference overlay [28]. However, such additions do not affect how the bars are plotted.

Y-axis Distortion. In comparison, another suite of alternative bar chart designs distort the Y axis in a bar chart thus affects how the length of the bars are determined and visually presented. Cut-off bars, scale break [11], and logarithmic scale are methods that were proposed to address problems classic linear bar charts may have in displaying disproportionately large values. These methods distort the Y-axis with a either non-linear scale or omitting a value range on the Y-axis. Cleveland discussed some of the techniques involving Y-axis distortion and called for experimentation for improving graphical communication in science [12]. Hlawatsch et al. noted that these distortions introduce a “lie factor” [33] when there is discrepancy of effect size between the data and its representation [20]. The distortions will result in misleading scenarios for quantitative comparisons, which is the most common task for classic linear bar charts [8]. More recently, Borgo et al. [6] proposed and evaluated an alternative bar chart design called Order of Magnitude Marker (OOMM) to facilitate the task of large magnitude number detection. The OOMM technique used a normalized scientific notation $A \times B^{10}$, with B determining the Y-axis scale and interval. The authors reported an empirical study (N=21) that demonstrates OOMM outperformed linear and log-scale bar charts design in identification and ration tasks. Note that the study did require participants to have basic knowledge of calculus and familiarity with concepts such as graphs and logarithmic scale. Therefore, we contend that OOMM and other techniques that distort the Y-axis may be better suited for communicating scientific data but may not be as effective for general audience [12, 11]. In contrast, the wrapped bar chart was originally invented for exhibitions and aimed to communicate discrepancy between values to general audience. For these reasons, our experiments evaluated the proposed wrapped bar chart technique against standard linear bar chart, without considering techniques that distort the Y-axis. Since the wrapped bar chart employ the same linear scale, we refer to standard linear bar chart as standard bar chart when reporting the experiments.

DESIGN OF DU BOIS WRAPPED BAR CHARTS

To develop wrapped bar charts inspired by Du Bois, we define two threshold variables that determine where and how to wrap a bar. First, we define a threshold $t1$ that determines where on the axis to wrap a bar for the first time. For a value exceeding $t1$, the bar representing this value will be wrapped. In terms of the wrapping direction, we decided to wrap down-to-up (first example in Figure 2). This design choice allows room for wrapping multiple times and ease of estimating the length of the wrapped portions.



Figure 3. Top: Standard bar chart. Middle: Wrapped bar chart at full length. Bottom: Wrapped bar chart at half length. Color choices for the bars and background are adopted from Du Bois' work in Figure 2.

Essentially, $t1$ limits the numerical axis by a value and wraps any bar that exceeds that value. In Figure 3, $t1$ is set to 1000. As a result, the largest bar with a value of 8500 has 8 full wraps (8 times $t1 = 1000$ equals 8000) with the tail reaching half of the numerical axis (equaling 500).

When designing the wrapped bars, we made multiple design decisions. First, in addition to a one-to-one ratio between the width of bars and the gaps between categories, we also determined on half of a bar-width gap for separating bars in the wrapped portion. Second, in our design, the value of a wrapped bar is the sum of only the vertical lines in a vertical wrapped bar charts and the wrapping portions in the horizontal axis do not contribute to the total value. Similarly, the vertical wrapping portions in a horizontal wrapped bar chart would not count toward the total length. Therefore, our wrapped bar chart maintains a linear relationship between values and vertical height of bars. Such design decision allows for easy estimate of the overall value of a wrapped bar. For example, a bar with a value of 5,500 in a wrapped bar chart with $t1 = 1000$ and $t2 = 1$, wraps a total of 5 times with 500 as the tail of the wrapped bar. These charts are developed using D3.js and an interactive version of the wrapped bar chart can be viewed at <https://wrapped-barchart.herokuapp.com/>.

Note that the wrapped bar chart prototype was first developed for Study 1 (shown in Figure 4). We improved the design of the wrapped bar chart based on feedback collected through Study 1. The changes were mostly cosmetic, including changing the bar and background color to resemble Du Bois' original work, and adjusting the bar width and gaps as detailed above. The final bar charts design used in Study 2 are illustrated in Figure 3 (top and middle charts).

STUDY 1: PILOT

Motivation and Hypotheses

Revisiting Du Bois' wrapped bar chart on African American students' enrolled courses (Figure 2-right), we can see that the number of students enrolled in industrial courses has orders

of magnitude larger than ones enrolled in business courses. From this example, we hypothesize that wrapped bar charts may bring two advantages when presenting datasets with high variance 1) the ability to maintain a fixed scale in a bar chart that allows for better visibility for bars with smaller values and 2) enabling more precise comparison/estimation of differences between larger and smaller values. However, we recognize that wrapping would discount the preattentiveness of the length of the bars [32], thus tasks of identifying the highest bar and estimating the value may take longer.

Based on these factors, we develop three hypotheses about wrapped bar charts in comparison to standard bar charts. Study 1 is conducted on two common types of tasks for a bar chart, namely **identification** and **ratio estimation** [6, 28, 13, 19, 30]. For these two types of tasks, we hypothesize that:

- **H1:** Participants will achieve higher accuracy with wrapped bar charts in identifying smallest values.
- **H2:** Participants will be more accurate with wrapped bar charts in estimating ratios involving smallest values.
- **H3:** Participants will take longer to complete of the tasks with wrapped bar charts, e.g., identifying the bar with the largest value, estimating the ratio of largest to smallest.

In addition to these hypotheses, we also want to measure the time performance of each participant per dataset to ensure that any gains in better task accuracy are not at the expense of additional time to complete. Therefore, we also report the time to complete the full set of tasks per individual per dataset.

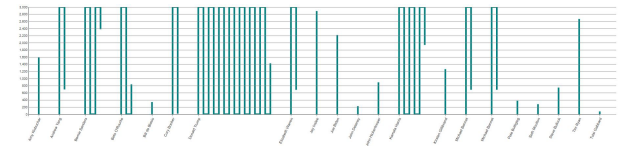
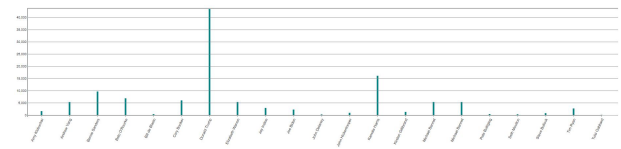
Experiment Design and Procedure

To compare participants' performance between wrapped bar charts and standard bar charts, we developed a web-based application for Study 1. The application recorded participants' responses and response time. The experiment design involves two factors: 2 datasets \times 2 chart types. We selected two real-world categorical datasets that exhibits very high variances between their largest and smallest categorical values (See Figure 4). The differences in these data values make the task of identifying and comparing lowest bars very difficult with standard bar charts. The first dataset is about number of Facebook ads by US presidential candidates for the 2020 election (values ranging from 78 to 43,000), while the second dataset is about number of resignation of the members of the US Congress by decade (data ranging from 1 to 122).

The experiment has a between-subject design as each participant sees either a wrapped or standard bar chart for each dataset. As shown in Figure 10, each participant viewed both datasets but with varying bar chart types by being randomly assigned to either group A or B. Bars in Study 1 were plotted with a fixed bar width and gap width calculated to evenly distribute the bars in the charting area. (see Figure 4). To catch participants that answered randomly, we added two simple bar charts (3-4 bars) before participants perform tasks on each dataset.

Each participant is asked to complete six tasks for each bar chart presented to them during the study:

Number of Facebook ads per Presidential Candidate



Number of Congressman Resignation per Decade

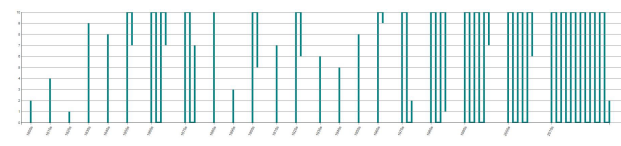
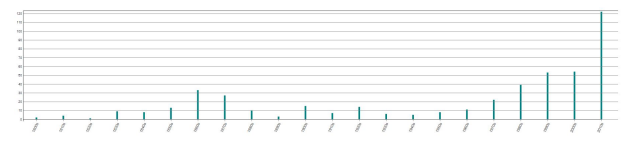


Figure 4. Two datasets used for Study 1. Top: Number of ad spending per presidential candidate, Bottom: Number of Congressional resignations per decade.

- **T1:** Identify the bar with the **largest** value.
- **T2:** Identify the bar with the **smallest** value.
- **T3:** How many times is the bar with the **largest** value to the bar with the **smallest** value?
- **T4:** Identify the bar with the **second largest** value.
- **T5:** Identify the bar with the **second smallest** value.
- **T6:** How many times is the value of the **second smallest** bar to the **smallest** bar?

For the identification tasks (**T1**, **T2**, **T4** and **T5**), the participants used a mouse click to select a bar and then click on the "submit" button. For the ratio estimation tasks (**T3** and **T6**), the participants enter a number and submit their answer.

Experiment Results

We deployed the study on Amazon Mechanical Turk. 98 participants completed the study with an average completion time of 5.5 minutes.³ Each participant was compensated \$0.50 for their time. In total, 56 participants were randomly assigned to Group A (standard bar first) and 42 to Group B (wrapped bar chart first).

We considered two different metrics to analyze the results. For **T1**, **T2**, **T4** and **T5** that resulted in categorical selection, we measured performance by user accuracy in identifying the

³One participant was dropped for not completing and a second participant was dropped for repeating one of the sections.

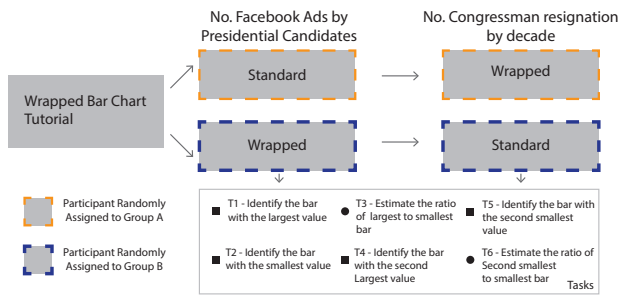


Figure 5. The general flow of Study 1. Each participant is assigned to one of the two (dataset*chart type) group after tutorial.

correct bar value (e.g., largest, smallest). We define user identification accuracy as the percentage of times they identify a small or large bar correctly and report the estimate differences in percentage points. For **T3** and **T6**, we used log absolute error similar to Cleveland and McGill [13].

To analyze and report our experiment data, we use a non-null hypothesis statistical testing (non-NHST) approach focusing on sample means and bootstrapped 95% confidence intervals (CIs) [16]. To enable comparison between our experiment results and other studies, we adopt the method by Talbot et al. to report estimates of simple effect sizes and associated CIs [30], i.e. differences between estimated outcomes of experiment conditions followed by 95% CIs in square brackets [10]. Narrow CIs compared to the estimated differences show strong evidence, while CIs that include zero imply more uncertainty about the sign of effect. As highlighted by Kim et al., reporting mean differences as simple effect sizes has limitations as variations around these averages are not considered [22]. To address such limitations, we also report Cohen's *d*, which is a measure of *standardized* effect size [14]. For consistency on the direction of reporting, we will report effect sizes and CIs of wrapped versus standard bar charts throughout this paper.

As an overview, Figure 6-left shows that participants achieved higher accuracy on the identification tasks by 10.07 [5.86, 14.28] ($d=0.32$) percentage points. Figure 6-right shows that participants achieved higher accuracy using wrapped bar charts across two datasets on ratio estimation tasks as indicated by estimated log absolute error difference of -2.84 [-3.91, -1.78] ($d=-0.52$).

Identification Accuracy Wrapped vs. Standard Bar Chart

Figure 7 shows the participants' accuracy on the identification tasks performed on both datasets. Consistent with **H1**, our results demonstrate the advantage of wrapped bar chart design in helping user identify bars with small values. Participants are more accurate with identifying the smallest value (**T2**) with wrapped bar chart design for both datasets. For the Facebook Ads dataset, participants were more accurate on average by 27.38 [13.69, 41.07] ($d=0.71$) percentage points using wrapped bar charts. Similarly, for the Congress dataset, participants were generally more accurate by 22.61 [8.92, 36.30] ($d=0.73$) percentage points. With the Facebook Ad dataset, participants are also more accurate at identifying the second largest value with wrapped bar chart design on average by

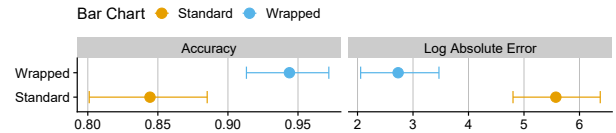


Figure 6. Study 1 mean and 95% individual-level bootstrapped CIs ($n = 98$) for identification accuracy (left) and log absolute error (right) for participants using wrapped bar chart vs. standard bar chart.

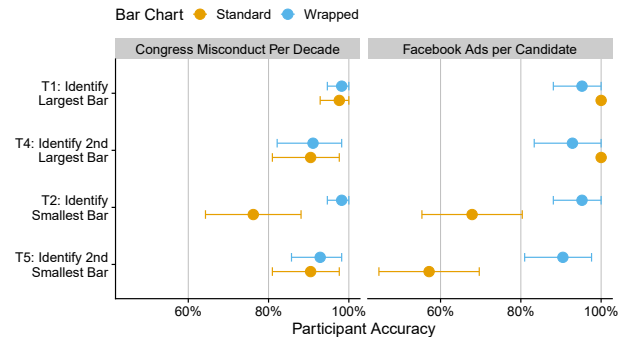


Figure 7. Study 1 mean accuracy and 95% bootstrapped CIs ($n = 98$) for identification task accuracy (T1, T4, T2, T5) using a wrapped vs. standard bar chart.

33.03 [17.26, 48.80] ($d=0.78$) percentage points. However, for the Congress dataset, we did not observe noticeable improvements for identifying second smallest 2.67 [-8.33, 13.69] ($d=0.08$). The analysis did not reveal noticeable differences in identifying largest and second largest bars in both datasets across the two conditions.

Ratio Estimation Accuracy Wrapped vs. Standard

Figure 8 shows the participant's performance on the ratio estimation tasks (**T3** and **T6**). The results for estimating the ratio of the largest to the smallest value (**T3**) is consistent with **H2** that participants with the wrapped bar chart performed much better. Specifically, the estimated mean difference in log absolute error is -4.30 [-5.36, -3.23] ($d=-1.52$) for the Facebook ads dataset and -3.01 [-4.38, -1.65] ($d=-0.8$) for the Congress dataset. Inconsistent with **H2**, no significant difference is observed between accuracy using wrapped and standard bar chart to estimate the ratio of second lowest to the lowest value.

Trial Completion Time

In Figure 9, we provide the mean completion time and 95% bootstrapped CIs per participant for each trial (i.e., dataset/chart combination with six tasks). We did not find evidence supporting **H3** as participants took nearly the same average time with wrapped bar chart as the standard bar chart for both datasets. However, one issue with only measuring the overall trial completion time is no visibility for the time taken for each individual task (e.g., largest bar identification, large / small ratio test). This issue is addressed in the second study.

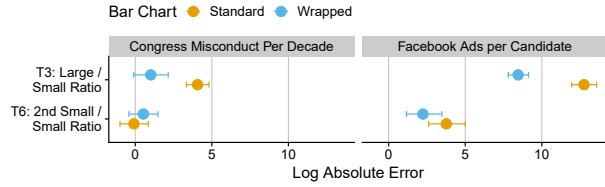


Figure 8. Study 1 mean and 95% bootstrapped CI's ($n = 98$) for log absolute error (T3, T6) of participants using wrapped bar chart vs. standard bar chart across tasks and datasets.

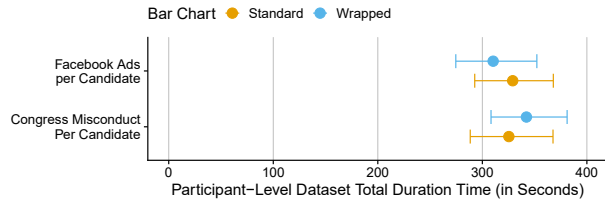


Figure 9. Study 1 mean completion time (in seconds) and 95% bootstrapped CI's ($n = 98$) per participant for all six tasks comparing wrapped bar chart vs. standard bar chart.

Discussion

The results from the pilot demonstrated that wrapped bar chart design can lead to improved accuracy for identification and ratio estimation task. However, what kind of dataset could benefit from being visualized in a wrapped bar chart remains unclear. As a result, we conducted another study detailed in the next section. The design of Study 2 is also informed by lessons learned through the first study design, implementation, and analysis. For example, the issue raised when analyzing the trial completion time motivates an improved data collection for Study 2 to better track time duration by individual task.

STUDY 2: INVESTIGATING WHEN TO USE WRAPPED BAR CHARTS

The primary goal of Study 2 is to investigate when it is advantageous to present data in a wrapped design instead of a standard bar chart. Study 2 is conducted to explicitly evaluate user performance on datasets of different characteristics.

Improved Study Design and Hypotheses

First, we highlight design changes from Study 2. Most importantly, to evaluate when to present data in a wrapped bar chart design, we leverage **theoretical metrics (normalized entropy and H-spread)** as heuristics to determine what data characteristics made them ideal candidates for wrapped bar charts.

Given our Study 1 results, we improved our study design and implementation in the following ways:

- Within-subjects study design to better measure the difference between standard versus wrapped bar charts.
- Remove identification and ratio test tasks of 2nd largest / smallest (T4, T5, and T6) to simplify study design and allow more repeated trials.

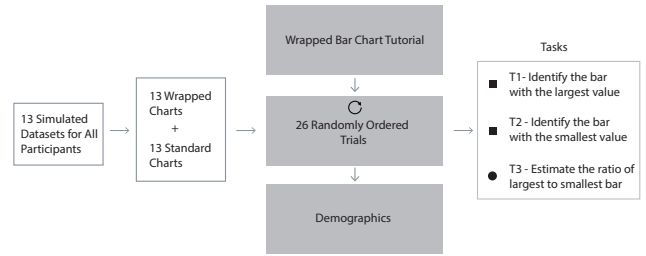


Figure 10. Study 2 Experiment Procedure

- Remove information about the dataset (such as number of Facebook ads in Study 1) to avoid distracting participants from focusing on the tasks.
- Enhance time tracking to measure task-level duration rather than only trial-level duration.

Study 2 is a within-subjects, repeated trials experiment that measures the efficacy of a wrapped bar chart relative to a standard bar chart. The goal of Study 2 is to provide guidance for visualization designers when to consider using wrapped bar chart based on the characteristics of the data to be visualized. In particular, we propose two data metrics to quantify what data characteristics render wrapped bar charts useful. Entropy serves as the primary metric for characterizing data distribution. However, as shown in the bar chart simulations (Fig. 11), each entropy bin can exhibit different top category concentration. We thus used H-spread as a secondary measure for how "far out" the highest bar is as defined by Tukey [34].

Primary Data Metric: Information entropy to characterize the concentration of values. In information theory, entropy measures the average (expected) amount of information from an event [18]. For a more certain event, there is less information contained in that event and implies a lower entropy. Conversely, an event that is equally likely across all possible outcomes (e.g., a uniform distribution) will have high entropy. Essentially, entropy is inversely related to concentration. We hypothesize that participants will perform better with wrapped bar charts for our tasks on discrete categorical datasets with low entropy because such datasets exhibit a disproportionate concentration of values within a few categories. Conversely, such performance gains with wrapped bar charts will diminish for discrete categorical with higher entropy as their values are more evenly spread across categories. We define entropy as:

$$\text{Entropy} = - \sum_{i=1}^N p_i * \log_2 p_i$$

where p_i is the percent of values for a category i and N is the number of categories. One issue with entropy is that it increases with the number of categories, limiting the comparison across categorical datasets with a different number of categories. Therefore, we normalized entropy by dividing it by the log (base 2) of the number of categories (N). This normal-

ization converts entropy into a range of values between 0 (all values concentrated into one bar) to 1 (uniform distribution).

$$\text{Normalized Entropy} = \text{Entropy} / \log_2 N$$

Secondary Data Metric: H-Spread to characterize disproportional values. In his book, *Exploratory Data Analysis*, John Tukey describes a method called ‘fences’ for identifying values that are “straying out far beyond the others” in a dataset [34]. His heuristic, which is widely used in boxplots to show disproportionate values, denotes the distance of a value away from the hinges (upper and lower quantiles) divided by the **H-spread** (the difference between the values of the hinges). He describes values larger than 1.5 times the H-Spread away from the hinges as ‘outside values’. Similarly, he names values that are 3 times the H-Spread outside the hinges as ‘far out’ values. Borrowing Tukey’s heuristic, we formally define the “H-Spread” metric for a categorical dataset with X values as:

$$\text{H-Spread} = (\max(X) - Q3(X)) / (Q3(X) - Q1(X))$$

Based on our results from Study 1 and our new theoretical metrics, we develop four hypotheses to test in Study 2.

- **H4:** Participants will achieve a **higher identification accuracy** for the smallest values with wrapped bar charts for datasets with **low normalized entropy**.
- **H5:** Participants will achieve a **higher identification accuracy** for the smallest values with wrapped bar charts for datasets with **high H-spread**.
- **H6:** Participants will have **better large-to-small accuracy** (i.e., lower log absolute error) in estimating ratios between largest and smallest bars with the wrapped bar charts than standard bar charts.
- **H7:** Participants will spend **more time** with wrapped bar charts when estimating largest-smallest bars ratio.

Data Simulation

To select our datasets for Study 2, we simulated 10,000 datasets based on 10,000 random draws from a fixed number of categories (e.g., 15 categories).⁴ Next, we categorized each distribution into four fixed ranges (bins) for both normalized entropy and H-spread, creating a 4×4 grid of 16 possible combinations. We then randomly drew one dataset per normalized entropy and H-Spread bin combination to use in our experiment. Given we used only a finite number of simulations, we were only able to generate datasets for 13 of the 16 because datasets with a low normalized entropy (less than 0.60) and high H-Spread (more than 4.50) did not occur in our 10,000 dataset simulation. Figure 11 provides the 13 sampled datasets used in Study 2 by each of the bin combinations. Although each dataset is visualized as a sorted bar chart in the data simulation app, the bars appear in random order in the simulated dataset for Study 2.

⁴Our simulator tool is a deployed R shiny app at https://ryanwesslen.shinyapps.io/wrapped_bar_sim/.

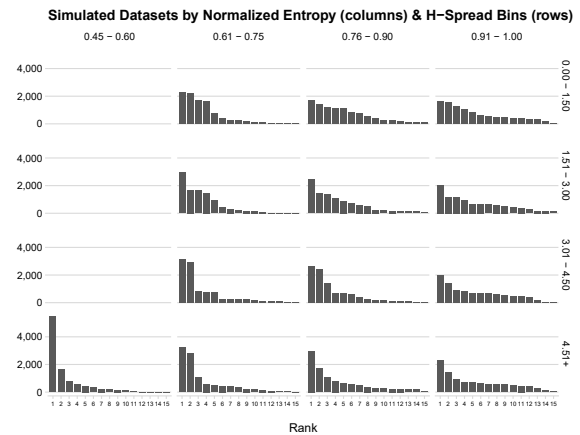


Figure 11. 13 simulated datasets used in Study 2 that were randomly sampled from each Normalized Entropy (Columns) and H-Spread (Rows) bin combinations for a 10,000 dataset simulation.

Experiment Design and Participants

For Study 2, we designed a within-subject Amazon Mechanical Turk experiment in which participants were provided both the standard and wrapped bar charts in random order for each simulated datasets. We improved the design of standard and wrapped bars to maintain a 1 to 1 ratio between bar width and gaps. Given 13 datasets, each participant completed 26 total trials consisting of three tasks: (**T1**) identify the largest bar, (**T2**) identify the smallest bar, and (**T3**) ratio estimation of the largest-to-the-smallest bar value. We randomized the order of the dataset-chart combinations (trials) to reduce related confounding factors.

203 participants completed our study, each receiving \$2.00 reward for completion. On average, participants took 29 minutes and 42 seconds. We evaluated individual level performance on identification tasks to identify participants who performed unreasonably. After investigating the number of correct answers on the easiest task (i.e., **T1**), we dropped 13 participants who incorrectly identified the largest value for both standard and wrapped bar charts for more than 10 of 26 datasets, leaving 190 participants for Study 2.

Experiment Results

For reporting our results, we used individual-level bootstrapped means since participants completed repeated trials across each experiment factor. To calculate, we first averaged performance on a participant basis per factor (e.g., Entropy bin) and then bootstrapped on user level performance. This enabled us to control for heterogeneity between participant performance, which is commonly found in crowdsourced visualization experiments [2]. Since Study 2 is a within-subjects study, we report paired Cohen’s d for within-subject paired samples [23].

First, consistent with **H4** and **H5**, we find participants in general had similar or better small bar identification accuracy with a wrapped bar chart compared to a standard bar chart. Figure 12 provides identification accuracy for Study 2 identification tasks (**T1** and **T2**) by normalized entropy and H-spread bin. The largest difference was in identifying the smallest value

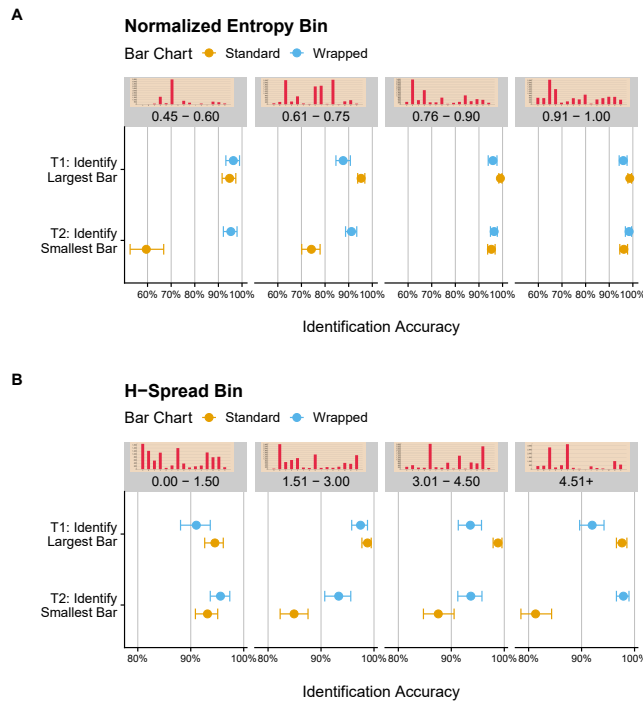


Figure 12. Study 2 large-small bar identification accuracy (T1 and T2) and 95% bootstrapped confidence intervals ($n = 190$) by Normalized Entropy bin (A) and H-Spread bin (B). We provide an example of a dataset from each bin as a visual cue.

bars in datasets with normalized entropy below 0.75. We observed a mean difference of 35.78 [28.42, 43.15] ($d=0.70$) percentage points for wrapped bar charts in datasets with normalized entropy between 0.45-0.6 normalized entropy bin; and a difference of 16.88 [12.31, 21.44] ($d=0.68$) percentage points for 0.6-0.75 normalized entropy. For the entropy bins of 0.75-0.9 and 0.9-1, the identification accuracy mean differences are small, respectively 1.06 [-1.01, 3.2] ($d=0.08$) and 2.01 [0.05, 4.26] ($d=0.2$).

In addition, we find that participants' had higher accuracy on small bar identification with wrapped bar charts than standard bar charts for datasets with H-spread larger than 1.5. We observe the largest effect with the highest H-spread category of 4.5+, with accuracy difference of 16.53 [13.36, 19.71] ($d=0.80$) percentage points. For H-spread of 3-4.5 we observe a difference of 5.92 [2.06, 9.78] ($d=0.28$) and for 1.5-3.0, a difference of 8.28 [4.64, 11.92] ($d=0.45$). We did not observe noticeable differences in datasets in H-Spread bin of 0-1.5.

Interestingly, we find wrapped bar charts yield worse accuracy in largest bar identification for datasets with high H-Spread. For datasets with H-Spreads of 3.00-4.5, participants were on average less accurate by -5.28 [-7.67, -2.89] ($d=-0.31$) and for datasets with H-Spreads of 4.5+, a similar worse accuracy of -5.65 [-8.15, -3.15] ($d=-0.34$) percentage points when using wrapped bar charts. For normalized entropy, we observe small or no effect within different value ranges. For normalized entropy values between 0.45 and 0.6, we did not observe significant differences ($d=0.05$) in largest bar identification accu-

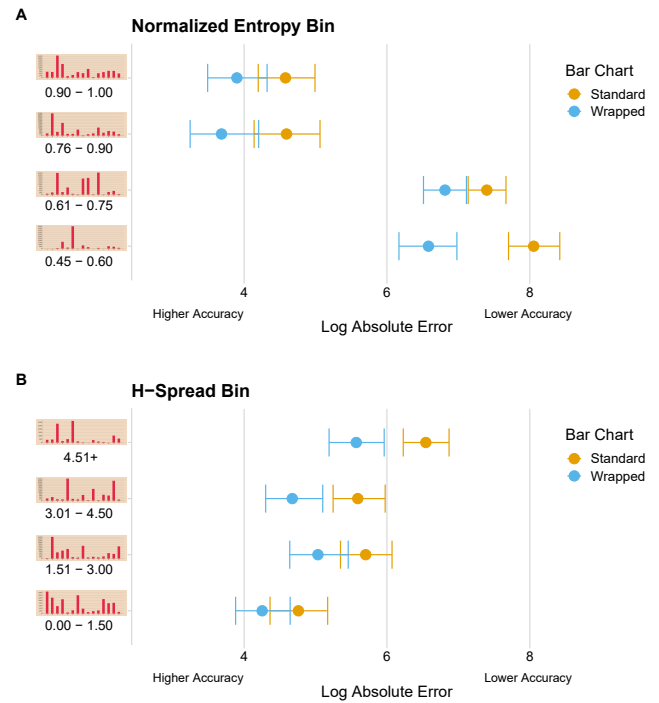


Figure 13. Study 2 the mean log absolute error and 95% bootstrapped confidence intervals ($n = 190$) for large-small ratio estimation (T3) by Normalized Entropy bin (A) and H-Spread bin (B). We provide an example of a dataset from each bin as a visual cue.

accuracy between wrapped and standard bar chart. For normalized entropy between 0.6-0.75, we observe on average a worse accuracy of -7.72 [-11.02, -4.42] ($d=-0.36$) percentage points with wrapped bar charts; for normalized entropy 0.75-0.9, we observe a mean difference of -3.26 [5.21, 1.31] ($d=-0.22$), and for 0.9-1.0 we observe a mean difference of -2.85 [-4.81, -0.89] ($d=-0.22$). We suspect that this may occur when multiple bars are wrapped, especially multiple times, leading to higher cognitive load in counting the number of wraps. We explore qualitative feedback to understand this hypothesis in our in-lab Study 3.

Second, we observe evidence for **H6** that participants consistently have better large/small ratio accuracy with wrapped bar charts compared to standard bar charts. Figure 13 provides the mean log absolute errors and 95% bootstrapped CI's by either normalized entropy or H-spread bins. For normalized entropy bins, we find the largest effect on datasets within the 0.45-0.6 range with a log absolute error mean difference of -1.47 [-2.01, -0.9] ($d=-0.62$). For normalized entropy within 0.6-0.75, we observe a mean difference of -0.58 [-0.97, -0.18] ($d=-0.54$). For normalized entropy of 0.75-0.9, we observe a mean difference of -0.89 [-1.54, -0.25] ($d=-0.57$). And we observe a mean difference of 0.67 [-1.26, -0.087] ($d=-0.56$) for normalized entropy values between 0.9-1.0.

For H-Spread bins, we find the highest mean difference of -0.97 [-1.49, -0.46] ($d=-0.69$) in the H-Spread 4.5+ datasets. We observe a similar mean difference of -0.91 [-1.47, -0.34] ($d=-0.55$) for datasets with H-Spread in the 3-4.5 range. We

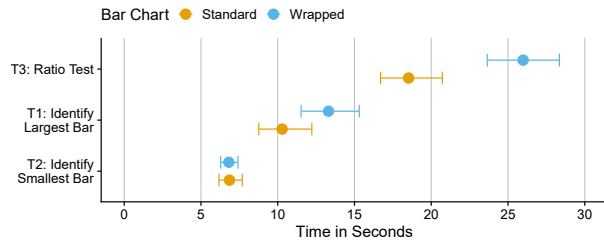


Figure 14. Study 2 mean completion time ($n = 190$) and 95% bootstrapped confidence intervals for each task.

observe a smaller effect of -0.67 $[-1.21, -0.14]$ ($d=-0.51$) for H-Spread 1.5-3 datasets. In summary, we observe that participants achieved higher accuracy in ration estimation in low normalized entropy and high H-spread datasets.

Third, we find some evidence in support of **H7** that on average participants tend to take 7.31 [4.98, 9.64] ($d=0.42$) seconds longer to complete the ratio task (**T3**) with a wrapped bar chart as compared to a standard bar chart. We find a smaller mean difference in time for the wrapped bar chart in identifying the largest bar **T1** of 3.04 [0.64, 5.44] ($d=0.18$) additional seconds. However, we find no difference in completion time for identification of the smallest bar **T2**. Figure 14 provides the mean time to complete each task along with 95% CIs. While we anticipated some increase in task time with the wrapped bar charts, overall, we do not find burdensome time effects that offset the benefits in identification and ratio accuracy.

STUDY 3: QUALITATIVE EVALUATION

Motivation and study design

After obtaining the quantitative results in Study 1&2, We design a qualitative study to gain a better understanding of users' individual experience with wrapped bar charts. We conducted in-lab focus group studies with a total of 24 participants. Our participants were graduate students enrolled in data science or architecture degree programs. None of these students had prior experience with wrapped bar charts. The procedure for Study 3 followed these steps:

1. Participants were given a brief overview of the study purpose (e.g., informed consent and instructions).
2. Participants completed an online study similar to Study 2. However, to control for time, participants were asked to complete tasks on 14 of the 26 datasets.
3. Participants completed an open-ended post-questionnaire individually to provide feedback on the study.
4. Participants take part in a focus group session on user experience, challenges, and suggested improvements.

The open-ended questions in the post questionnaire asked participants for their strategies on calculating ratios and their feedback on the usefulness of wrapped bar charts. After all participants in a session completed the study, we held group discussion about wrapped bars to further understand participants' feedback. We collected participants' written comments,

transcribed audio recordings from the open discussion sessions. Two coders independently thematically analyzed users comments and feedback. We identified four themes from participants' written comments and discussions:

Benefits of wrapped bar charts: The most prevalent themes from both focus group studies include the intuitiveness of the wrapped bar chart design, when they could be beneficial, and participants' strategies to interpret wrapped bars. Out of 24 participants, 21 mentioned they believe wrapped bar charts are useful for specific cases when some values in a dataset are disproportionate. One participant's comment summarizes this theme: "[*Wrapped bar charts are useful*] when there is an obvious disparate effect from one attribute over the others. If the axis scales were calibrated effectively, wrapped bar charts could more concisely convey differences in orders of magnitude". Another participant understood the benefit of wrapped bars but highlighted the importance of the tasks that chart designer want to explicitly support: "[*Wrapped bar charts are useful*] when it is hard to read minimums and when the minimums are crucial."

Usability of wrapped bar charts: Participants highlighted two important points on the usability of wrapped bar charts. First, they understood that wrapped bars were useful for specific cases when largest values render a number of small values illegible. However, 16 participants argued that at some point the task of estimating values of the wrapped bars can be cumbersome, especially when there are many wraps. One participant commented: "*wrapped chart is easier if there are only couple of wraps. but if number of wraps increase, it gets harder*". A major source of annoyance for five participants was when they had to estimate the tail end value of wrapped bars stems from the opposite direction of the y axis (for the wrapper portion going top to bottom), as one participant mentioned: "... *You need to do additional subtractions when the bar is coming downwards*".

Strategies for reading wrapped bar charts: Another emerging theme in our discussions with participants was the different strategies they took to read wrapped bars, and calculate ratios between largest and smallest values. Majority of users simply tried to estimate the value of largest and smallest, and do a division task as described by two participants: "*I assume a number for minimum between the range. Then, another number for the max. After that, I just divide and make the number round if possible.*" and "*count the tall bar then the short bar and divide. for the tall bars I counted the number of full bars and multiplied by the max number then added the end bar up.*" However, three participants came up with another novel strategy to read wrapped bars: "*I started estimating values of each of the largest and shortest columns [bars] and then divided them. But after a few iterations I found a better method. I estimated if stacking the shortest [bar] on itself X number of times would fill up to the first grid line and then counted how many lines the longest [bar] took up.*"

Suggestions for improving wrapped bar charts: In the focus-group sessions, participants also discussed how to improve the design of wrapped bar charts. One suggestion was including an 'inverse axis' to the right side of a wrapped bar

chart to help estimate the value of the tail end of wrapped bars: “Use two y scales on both left and right side of the chart but in a reverse way”. Moreover, participants had multiple suggestions on ways of simplify the estimation of value of largest bars by including count of wraps, using colors, and increasing gaps between wrapped bars.

DISCUSSION AND FUTURE WORK

In our two online experiments, we find specific cases (low entropy and high H-Spread) where users gain accuracy in identification and ratio tasks while using wrapped bar charts. However, we also identify subtleties and drawbacks in using wrapped bars in these specific cases. First, both quantitative and qualitative findings including participants’ higher time spent on identifying the largest bar and mentions of annoyance with too many wraps provide evidence that wrapped bar charts are less preattentive than standard linear bar charts. Our participants recognized that for estimating values of wrapped bars, they need to go through a mathematical process of addition or multiplication of values. Based on the feedback in the focus-group study, we believe that this effect might be negligible for a small number of wraps. Adding information such as providing a label for the number of wraps is needed for larger number of wraps. To understand the trade-off between cognitive load and accuracy in wrapped bar charts, we plan to conduct a future experiment studying the optimum threshold for wrapping bar charts and the limit to the number of times a bar can be wrapped before it becomes ineffective.

To understand the source of error w.r.t the ratio estimation task, we further explored whether identification accuracy impacts users’ ratio estimation accuracy. Out of 4,949 valid trial responses collected in Study 2, 12.5% (617) inaccurately identified either the largest or smallest bar. We ran a mixed effects model on only the trials with correct identification responses and found similar effects on log absolute error, with both means decreasing by about 0.1. This exploration suggests an improvement for future study design, i.e. to explicitly highlight the correct highest and smallest bars for the ratio estimation task.

In the future, there are several design improvements we plan to make and evaluate with wrapped bar charts in order to reduce the cognitive load required to interpret wrapped bars. First, as mentioned in the focus-group studies, we plan to introduce an inverse y-axis to the right side of a horizontal bar chart to help read the tail end of wrapped bars (when going from top to bottom). Moreover, we also plan to experiment different gap sizes between wrap bar charts to help with the task of counting the number of wraps. Third, we plan to study the effects of changing the second threshold of wrapped bars (see Figure 3 bottom) on participant performance. Finally, one solution to the trade-off between cognitive load and benefits of wrapped bars could be to introduce a new interaction technique for interactive bar charts. ‘Wrapping’ would allow users to quickly move from standard bar charts to wrapped bar charts and observe the values of small bars. These design solutions require a comprehensive and systematic study to estimate and understand the benefits of these changes.

From an application perspective, there are multiple extensions we can test with wrapped bar charts. First, in order to better understand the potentials of wrapped bar charts, we need to collect more data from users including mouse positions, user confidence level, and user satisfaction. Second, we need to compare wrapped bar charts with other types that are built to deal with datasets with values of high variance, such as broken axis and logarithmic scale. Third, during our development of wrapped bar charts, we developed a new kind of visual interaction by allowing users to interactively change the thresholds of wrapped bar charts (t_1 & t_2) and explore high and low data points. In a future study, we plan to thoroughly study this interaction technique for conducting similar visual tasks and learn about user preference. Last, we suspect that zoomed-out wrapped bar charts could work well with many more categories ($N > 50$), like a power law distribution (e.g., Zipf’s law for word counts) that is carefully ordered (e.g., by word-topic theme).

Finally, as a general suggestion for visualization designers to employ wrapped bar chart design within their work we provide the following recommendation resulting from our quantitative and qualitative evaluations: Using entropy and H-Spread as heuristics to measure a dataset and determine whether it is beneficial to use wrapped bar charts. We recommend using wrapped bar charts to visualize datasets with normalized entropy less than 0.75 as we see performance gain in identification and ratio tasks (Figures 12 and 13). Moreover, in line with Tukey’s suggestion, when the H-Spread threshold is larger than 4.5 that includes far out values), wrapped bar charts could be a good design option.

CONCLUSION

In this paper, we designed, developed, and evaluated wrapped bar charts originally introduced by W.E.B. Du Bois to highlight the benefits of his visual innovation added to the standard linear bar charts presentation. We developed an implementation of wrapped bar charts for the web. We hypothesized that wrapped bar charts can outperform standard bar charts for datasets with disproportionate values. Using a web-based interface we conducted two online experiments. We found that for cases where data values are disproportionate as measured by normalized entropy and H-spread, wrapped bar charts allow participants to achieve higher accuracy on tasks of identification and ratio estimation but sometimes at the expense of more time spent and potentially more cognitive load. Finally, resulting from our focus group study, we develop a list of potential design improvements for such charts. Our findings, can serve as guidelines for visualizing datasets with disproportionate values using wrapped bar charts.

REFERENCES

- [1] 2016. W. E. B. Du Bois’ Hand-Drawn Infographics of African-American Life (1900) The Public Domain Review. <https://publicdomainreview.org/collections/w-e-b-du-bois-hand-drawn-infographics-of-african-american-life-1900>. (2016). (Accessed on 06/12/2019).
- [2] Alfie Abdul-Rahman, Karl J Proctor, Brian Duffy, and Min Chen. 2014. Repeated measures design in

- crowdsourcing-based experiments for visualization. In *Proceedings of the Fifth Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization*. ACM, 95–102.
- [3] Witney Battle-Baptiste, Britt Rusert, and others. 2018. *WEB Du Bois's Data Portraits: Visualizing Black America*. Chronicle Books.
- [4] James R Beniger and Dorothy L Robyn. 1978. Quantitative graphics in statistics: A brief history. *The American Statistician* 32, 1 (1978), 1–11.
- [5] Rita Borgo, Alfie Abdul-Rahman, Farhan Mohamed, Philip W Grant, Irene Reppa, Luciano Floridi, and Min Chen. 2012. An empirical study on using visual embellishments in visualization. *IEEE Transactions on Visualization and Computer Graphics* 18 (2012), 2759–2768.
- [6] Rita Borgo, Joel Dearden, and Mark W Jones. 2014. Order of magnitude markers: An empirical study on large magnitude number detection. *IEEE transactions on visualization and computer graphics* 20, 12 (2014), 2261–2270.
- [7] Eliot Brown. 2019. This Is Not the Way Everybody Behaves. How Adam Neumann's Over-the-Top Style Built WeWork. <https://www.wsj.com/articles/this-is-not-the-way-everybody-behaves-how-adam-neumanns-over-the-top-style-built-wework-11568823827>. (Sep 2019). (Accessed on 09/19/2019).
- [8] John Bryan. 1995. Seven types of distortion: A taxonomy of manipulative techniques used in charts and graphs. *Journal of technical writing and communication* 25, 2 (1995), 127–179.
- [9] Aaron Bycoffe, Sarah Frostenson, and Julia Wolfe. 2019. A Final Look At Who Won The Third Democratic Debate. <https://projects.fivethirtyeight.com/democratic-debate-september-poll/>. (Sep 2019). (Accessed on 09/16/2019).
- [10] Guillaume Calmettes, Gordon B Drummond, and Sarah L Vowler. 2012. Making do with what we have: use your bootstraps. *Advances in physiology education* 36, 3 (2012), 177–180.
- [11] Clark R Carvalho and Michael D McMillan. 1992. *Graphic representation in managerial decision making: The effect of scale break on the dependent axis*. Technical Report. AIR FORCE INST OF TECH WRIGHT-PATTERSON AFB OH.
- [12] William S Cleveland. 1984. Graphical methods for data presentation: Full scale breaks, dot charts, and multibased logging. *The American Statistician* 38, 4 (1984), 270–280.
- [13] William S Cleveland and Robert McGill. 1984. Graphical perception: Theory, experimentation, and application to the development of graphical methods. *Journal of the American statistical association* 79, 387 (1984), 531–554.
- [14] Jacob Cohen. 1988. *Statistical power analysis for the social sciences*. (1988).
- [15] Geoff Der and Brian S Everitt. 2014. *A handbook of statistical graphics using SAS ODS*. Chapman and Hall/CRC.
- [16] Pierre Dragicevic. 2016. Fair statistical communication in HCI. In *Modern Statistical Methods for HCI*. Springer, 291–330.
- [17] Stephen Few. 2013. Wrapping Graphs to Extend Their Limits. https://www.perceptualedge.com/articles/visual_business_intelligence/wrapping_graphs_to_extend_their_limits.pdf. (2013).
- [18] Robert M Gray. 2011. *Entropy and information theory*. Springer Science & Business Media.
- [19] Jeffrey Heer and Michael Bostock. 2010. Crowdsourcing graphical perception: using mechanical turk to assess visualization design. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 203–212.
- [20] Marcel Hlawatsch, Filip Sadlo, Michael Burch, and Daniel Weiskopf. 2013. Scale-Stack Bar Charts. In *Computer Graphics Forum*, Vol. 32. Wiley Online Library, 181–190.
- [21] Thomas Kaplan and Sarah Almkhatar. 2019. How Trump Is Outspending Every 2020 Democrat on Facebook. <https://www.nytimes.com/interactive/2019/05/21/us/politics/trump-2020-facebook-ads.html>. (May 2019). (Accessed on 09/13/2019).
- [22] Yea-Seul Kim, Jake M. Hofman, and Daniel G. Goldstein. Effectively Communicating Effect Sizes. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. Association for Computing Machinery, New York, NY, USA.
- [23] Daniël Lakens. 2013. Calculating and reporting effect sizes to facilitate cumulative science: a practical primer for t-tests and ANOVAs. *Frontiers in psychology* 4 (2013), 863.
- [24] Pranathi Mylavarapu, Adil Yalcin, Xan Gregg, and Niklas Elmqvist. 2019. Ranked-List Visualization: A Graphical Perception Study. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA, Article Paper 192, 12 pages. DOI: <http://dx.doi.org/10.1145/3290605.3300422>
- [25] Drew Skau, Lane Harrison, and Robert Kosara. 2015. An evaluation of the impact of visual embellishments in bar charts. In *Computer Graphics Forum*, Vol. 34. Wiley Online Library, 221–230.
- [26] Ian Spence. 2006. William Playfair and the psychology of graphs. In *Proceedings of the American Statistical Association, Section on Statistical Graphics*. 2426–2436.

- [27] Ian Spence and Stephan Lewandowsky. 1991. Displaying proportions and percentages. *Applied Cognitive Psychology* 5, 1 (1991), 61–77.
- [28] Arjun Srinivasan, Matthew Brehmer, Bongshin Lee, and Steven M Drucker. 2018. What’s the Difference?: Evaluating Variations of Multi-Series Bar Charts for Visual Comparison Tasks. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 304.
- [29] Ben Starr. 2014. How to Design Bar Charts. <https://visage.co/data-visualization-101-bar-charts/>. (2014). (Accessed on 06/13/2019).
- [30] Justin Talbot, Vidya Setlur, and Anushka Anand. 2014. Four experiments on the perception of bar charts. *IEEE transactions on visualization and computer graphics* 20, 12 (2014), 2152–2160.
- [31] AnyChart Team. 2015. First Ever Bar Charts Created in 14th Century. <https://www.anychart.com/blog/2015/06/02/first-bar-chart-in-history/>. (June 2015). (Accessed on 06/13/2019).
- [32] Anne Treisman and Janet Souther. 1985. Search asymmetry: A diagnostic for preattentive processing of separable features. *Journal of Experimental Psychology: General* 114, 3 (1985), 285.
- [33] ER Tufte and P Graves-Morris. 1983. Vol. 2. *The Visual Display of Quantitative Information*. CT: Graphics press Cheshire (1983).
- [34] John W Tukey. 1976. Exploratory data analysis. *Massachusetts: Addison-Wesley* (1976).
- [35] Mingqian Zhao, Huamin Qu, and Michael Sedlmair. 2019. Neighborhood Perception in Bar Charts. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, Article 232, 12 pages.