Applying Racial Equity Awareness in Data Visualization

Jonathan Schwabish*
The Urban Institute

Alice Feng**
The Urban Institute

ABSTRACT

A data visualization style guide does for graphs what the *Chicago Manual of Style* does for English grammar: it defines the components of a graph and their proper, consistent use. At the Urban Institute, a nonprofit research institution based in Washington, DC, our data visualization style guide defines these styles for our research and communications staff. In the process of revising and expanding the style guide, we are taking a more diverse, equitable, and inclusive (DEI) perspective to our research, data, and visualizations. To date, our approach has been to create a set of recommendations and issues to consider, rather than a set of rules that researchers must follow. In this paper, we discuss eight techniques that data visualization producers should consider when creating visuals with this DEI approach. As with our existing style guide, we consider this a first step in our thought process of creating graphs and producing content through a DEI lens.

Keywords: Diversity, equity, inclusion, racial equity, accessibility.

Index Terms: Accessibility, codes of ethics, user characteristics.

1 INTRODUCTION

The Urban Institute is a nonprofit research institution based in Washington, DC. With more than 500 researchers, staff members, and communications experts, Urban's mission is to open minds, shape decisions, and offer solutions. To advance that mission, we must think intentionally about how we can learn from and speak to audiences that reflect the rich diversity of America's communities, and we must foster a culture where employees from different backgrounds and perspectives enjoy mutual respect, inclusivity, and collegiality.

For the past few months, we have been working on revising and updating Urban's Data Visualization Style Guide [16]. In addition to reorganizing the guide and updating various style decisions, we have also been thinking critically about how we communicate data and information about the groups we study. Just as Urban has carefully considered the words we use in our written reports and platforms, we should be equally careful in how we visually present data to our readers, users, and audiences and what words we should use in and around those visuals. (We are also working on a separate section about accessibility in data visualization, which we will share at a later date.)

Although more people are thinking and writing about these issues, there hasn't been much agreement around best practices for taking an equity lens to data visualization, especially as it applies

---

* Jonathan Schwabish is a senior fellow at the Urban Institute. E-mail: jschwabish@urban.org.
** Alice Feng is a data visualization developer at the Urban Institute. E-mail: afeng@urban.org.

to setting standards for entire organizations. As best we can, we have been reading a variety of posts and papers (a short list can be found below) and discussing ways we can develop a more diverse, equitable, and inclusive approach to presenting and visualizing data. We view this effort as just the beginning of our process and anticipate growing and expanding our work as we learn more and receive feedback from colleagues, partners, and readers.

To that end, we have identified eight areas in which researchers, analysts, and anyone working with data can be more inclusive in how they present their data.

2 USING LANGUAGE WITH A RACIAL EQUITY AWARENESS

Titles, text, and labels are among the first things readers scan [3] when encountering a chart, so they present an important opportunity to apply racial equity awareness thinking. Urban researchers often rely on our organization's guides for accessible and inclusive language for labeling and annotating graphs. Consistent with Urban style, we want our researchers to label their data using people-first language, such as "people with disabilities" rather than "disabled people." Terms should also refer to people and not strictly to their skin color; for example, "Black people" not "Blacks."

As just one example, in a June 2020 project [1], a series of bivariate choropleth maps showed the relationship between race and poverty. In the original visualization, the labels along the legend were "More Poverty" and "More Black." That language is not inclusive of different groups: poverty refers to an experience not a static description and "More Black" references skin color, not people. A more inclusive way to label the legend might be "Larger proportion of people experiencing poverty" and "Larger Black Population" (the author of the visualization did later change "More Black" to "Larger Black Population").
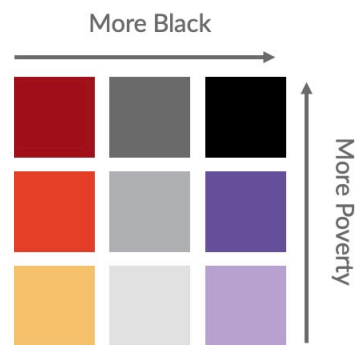

Figure 1: Legend from the project. For purposes of image clarity, this legend was remade from the original.

The terms and phrases we use continue to change. In writing about terminology around and about people with disabilities, Nicholas Steenhout [14] writes, "Disability language is never straightforward. It's always nuanced. It always evolves." That sentiment can extend to any underrepresented group, so as researchers, developers, and designers, we need to monitor the current lexicon and reflect the experiences of the people we study

and the people we communicate with. Graph creators may want to consider including a footnote or endnote explaining why a particular term was used. We also encourage researchers to talk to their target audiences to give them the option to self-identify their preferred terminology to determine what language is most appropriate for them.

3 ORDERING DATA LABELS IN A PURPOSEFUL WAY

Many of our graphs and tables include different demographic groups, ordered simply as they appear in the raw data. There is likely little thought given to *how* estimates in tables or bars in graphs or small multiples are ordered, such as "White," "Black," "Hispanic," "Other." As graph producers, however, we should take a more active role in choosing how to order and present data values for different groups. Which group we choose to show as the first row in a table or the first bar in a graph can affect how readers perceive the relationship or hierarchy between groups; always starting with "White" or "Men" can make these groups appear as the default against which other groups should be compared, suggesting they're the most important populations. How we choose to order may also reflect who we view as the intended audience for our visualizations. Starting with "White" or "Men" can make it seem as though those are the most important groups we are trying to communicate with. Urban does not have a universal rule that applies to all visuals, but a few issues are worth considering:

- Does your study focus on a particular community? If it does, that group should be presented first.
- Is there a particular argument or story you are trying to tell? If so, the order or presentation of results should reflect that argument.
- Is there a quantitative relationship that can guide how the groups are ordered? Can they be sorted alphabetically or by population size, sample size (weighted or unweighted), or magnitude or effect of the results?

4 CONSIDERING THE MISSING GROUPS

It is also important to acknowledge who is and is not included in our data and charts. Many charts on race and ethnicity only show white people, Black people, and Latinx people but not smaller racial or ethnic groups. Often this is because of data limitations, in particular, small sample sizes. But even in those cases, how can researchers be more proactive to help organizations conducting surveys be more inclusive? How can we communicate to those organizations to help them conduct better surveys? Just because it may be harder to obtain data about certain groups doesn't mean we shouldn't still try to better understand their lives.

Similarly, charts showing breakdowns by gender often neglect nonbinary and transgender people by only presenting males and females. Again, many major surveys, especially at the federal level, do not include these groups as response options, but we should be asking ourselves how and to what extent we should note omitted groups and when. Should researchers assume readers will recognize that certain data sets treat gender as binary? How can we anticipate that?

Another topic we've been considering is how to avoid lumping groups with few observations into the seemingly innocuous "Other" category. In some cases, the "Other" category may be necessary to achieve sufficient sample size for statistical analyses—in which case, we should ask ourselves if it is appropriate to lump these groups together (do they show the same trends in the data?) and be sure to clearly define in the notes section of the chart who is included in this category—but in other cases, we may use the "Other" category as a shortcut to make the presentation and related writing easier. Should we start considering words besides "Other," which can have an exclusionary connotation to label these groups? Even though it might take more

words, maybe we should use the appropriate name for each group in our tables, charts, and text.

In cases where data were collected about a specific group but that group was not presented separately in the chart, it may be worth listing in the chart's notes all of racial and ethnic groups included in the original dataset. This can both inform readers that data do exist for these smaller groups (even though they may be limited) and offer transparency about the chart maker's decision about groups they included in their visualization.

5 USING COLORS WITH A RACIAL EQUITY AWARENESS

Urban's color palette is consistent for people with certain color vision deficiencies, and the contrast between those colors and white/black text meet basic accessibility guidelines [18]. Urban does not use color palettes that reinforce gender or racial stereotypes, such as baby pink and baby blue to represent women and men. We have not set a specific standard for which Urban color refers to which gender group or racial group, although we certainly recommend that researchers avoid using colors associated with skin tones (or worse, racial stereotypes).

As an example, the Massachusetts Institute of Technology Office of the Provost's "Diversity Dashboard" enables users to explore demographic characteristics of the school's students, faculty, and staff. The June 2020 version of the dashboard used three distinct hues to represent nine racial and ethnic groups. Five of the distinct race and ethnicity groups were represented with red; a similar lighter pink color for the "Two or more races" category; two shades of gray for the "International" and "Unknown" groups; and blue for the "White" group.
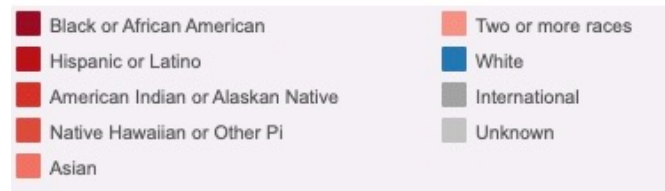

Figure 2: Original legend from the "Diversity Dashboard."

There are multiple problems with this design. First, the red color ramp for the five groups on the left side of the legend others [17] those groups and creates a visual divide that seems to pit nonwhite students against white students. Secondly, the gray colors tend to fade to the background [13], which diminishes international students and students whose race or ethnicity is unknown. All of this creates an effect where the "White" category moves to the foreground and is highlighted, as if that is the most important group and norm to which all other groups should be compared. An update to the dashboard in July 2020 partly remedied the color issues by using more hues and creating larger breaks between the now-blue color ramp.

Figure 3: Updated legend from the "Diversity Dashboard."

In general, as data visualization producers, we need to be aware of how our use of colors, words, and categorizations can perpetuate or exacerbate inequities.

6 USING ICONS AND SHAPES WITH A RACIAL EQUITY AWARENESS

We want to be careful and thoughtful when using icons in any data visualization. When showing groups of people, we should consider a mix of genders, races, and ethnicities. We don't often use icons at

Urban, but when we do, we should consider to whom we are presenting our results and how our icons might be perceived. We need to be conscious of how certain icons may not correspond to the content, such as an icon of a baby in a chart about child mortality.

Mis- or underrepresentation of certain groups in imagery and iconography also fails to take a racial or gender equity awareness perspective toward our data visualization work. A 2018 study by the Pew Research Center found that "men are overrepresented in online image search results across a majority of jobs examined" and that "women appear lower than men in such search results for many jobs." This disparity continues to this day; the image on top shows Google search results for the phrase "nurse icon," and the image on the bottom shows results for the phrase "boss icon." Notice how nearly all the images on the left are similar to images we might imagine as feminine and the images on the right we might traditionally consider masculine.

Figure 4: Google image search results for "nurse icon" (top) and "boss icon" (bottom).

Data visualization creators should ensure a variety of races and genders are depicted [9] when using icons and avoid icons that make inappropriate depictions of people or communities or reinforce stereotypes.

7 DEMONSTRATING EMPATHY

One of the big challenges in visualizing data, and quantitative research in general, is the ability to help readers connect with the content. Standard graphs like bar charts, line charts, and pie charts, while informative, can abstract from the content and people being represented. Taking an empathetic view of the reader's needs as they read or perceive information is one step to better data communication [12]. This kind of empathy is often couched in terms of producing specific graphs that meet the needs and expertise of our readers. But viewing empathy through a DEI lens would mean considering how the specific lived experiences and perspectives of our readers (not to mention the actual study populations) will perceive the information. As journalist Kim Bui [5] wrote in 2019, "approaching stories—and people—with more

empathy creates better relationships with marginalised communities, builds trust and increases diverse coverage."

Empathy in data visualization also extends to when we use abstract shapes versus shapes that resemble people or communities, which can evoke a little more humanity. Graphics that specifically represent people—the anthropomorphizing of data graphics, or, as Jeremy Boy and colleagues [4] refer to it, anthropographics—is sometimes seen as a way to evoke empathy (though Boy and his coauthors do not find this to be the case; see also Groeger [7]). The idea of helping readers understand the "near and far"[8]—a wide lens of overall metrics or data paired with individual- or group-level data—may be a way to help readers connect with content [10]. The guiding principle here is to put people first and help the reader better understand and recognize the people behind the data. "One possibility," web developer Jacob Harris wrote, "is that if your data is about people, make it extremely clear who they are or were."

Another way to think about empathy in data visualization is whether particular chart types lend themselves to more of a human connection than others. For example, bar charts represent each data value with a single rectangle. On the other hand, unit charts and waffle charts visualize each data point using a collection of smaller, individual shapes.

Figure 5: Example of a bar (top) versus unit chart (bottom).

The former approach could be seen as more abstract, collapsing all of the people reflected in that data point into just one shape, while the latter chart type might offer more of an opportunity to connect with the subject by reminding readers of the number of people represented through the use of multiple individual shapes, particularly if each dot represents one person.

Connecting directly with people and communities and trying to better understand their lived experience can help content producers create visualizations and tell stories that better reflect the true experiences of different people. This connection may come more naturally for journalists than for data scientists, but we believe this is a key dimension for researchers to explore to help them put their work into the hands of policymakers, stakeholders, community members and changemakers who can use it to affect change. Furthermore, inclusive and thoughtful data visualization that respectfully reflects the people we are studying can also help us build trust with those communities.

8 QUESTIONING DEFAULT VISUALIZATION APPROACHES

It is often reflexive when presenting data disaggregated by race or ethnicity to plot all of the groups on the same chart. However, as Pieta Blakely [2] writes, doing so "encourages [the reader] to compare each of the groups to the highest performing groups" which can lead to a "deficit-based perspective" that focuses attention on what low-performing groups are lacking compared with the high performers. Instead, Blakely proposes plotting each race on its own chart as a set of small multiples, which can better encourage readers to think about the specific needs and challenges facing each group.

This was the approach we took with Urban's Tracking COVID-19's Effects by Race and Ethnicity [15] data tool, which seeks to present the disparate impact of the COVID-19 pandemic by race and ethnicity for various metrics. Initially, line charts were proposed that displayed all of the races in the same chart.
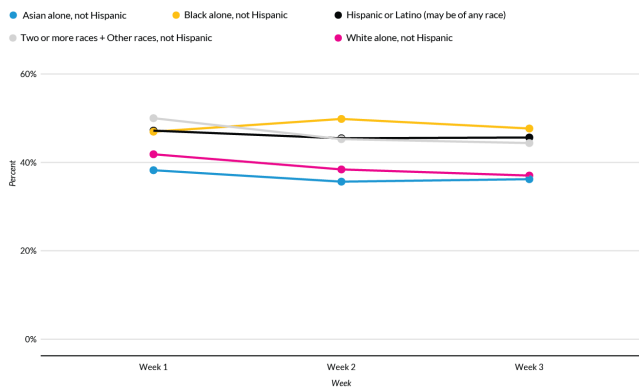
Figure 6: Original sketch for Tracking COVID-19's Effects by Race and Ethnicity.

But as our work on the project progressed, we applied a racial equity awareness lens to how we were visualizing the data and began to question our initial approach. For one, putting all of the races on the same chart had the effect of focusing attention on only the hardest-hit groups, rather than on how members from all races and ethnicities have been negatively affected by the pandemic. Another issue was that plotting all of the races on one graph made it seem like the least-affected group was the goal the other races should seek to achieve. But why should that be the case? Upon further questioning, we decided a better comparison would show the state- or metro-wide average. This approach had a couple benefits. First, it offered a consistent standard that all groups could be measured against. Secondly, it meant we did not promote any particular racial or ethnic group as being the default group against which everyone else must be compared.
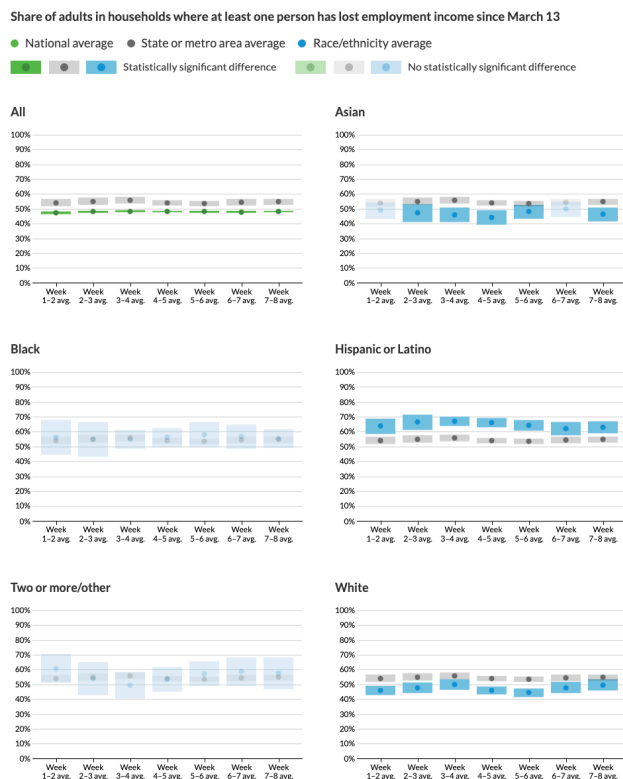

Figure 7: Final design for Tracking COVID-19's Effects by Race and Ethnicity.

Our final design presents the data using a set of small multiples faceted by each racial and ethnic group with the relevant local average included as a benchmark. The estimates are shown as points surrounded by confidence interval blocks because of statistical and data collection concerns. The distance between the grey and green or grey and blue blocks reflects how much better or worse the group is doing compared with the average, and the change in that distance week to week indicates whether the disparity is worsening or improving over time. In weeks when the difference was not statistically significant, the estimates are shown using a lower opacity to reflect our uncertainty about the magnitude of the gap. We also made a conscious decision to sort the groups alphabetically rather than by data value. Because users can select different metrics in the final project, sorting alphabetically maintains the same order across all views rather than having the graphs shift around.

9 ENGAGING OR REFLECTING LIVED EXPERIENCE

Ideally, chart makers would not just implement our guidelines above but would also reach out to members of the communities being visualized and ask for their feedback and advice. Are we using language consistent with how they refer to themselves and others? What have we missed in our visuals that are inconsistent with a DEI framework? How can we take a more empathetic approach to creating data visualizations that accurately and respectfully account for other people's lived experiences? It's important to remember that data are a reflection of the lives of real people, not just a sterile abstraction.

10 LOOKING FORWARD

As we continue to revise our data visualization style guide, we've noticed few, if any, style guides from other organizations [11] mention accessibility, inclusion, or diversity. Our approach has been to create a set of recommendations and issues to consider rather than a set of rules that researchers must follow. Ultimately, we hope researchers will be thoughtful and deliberate about their design choices [6], not just relying on software defaults or the status quo. If you or your organization has gone about setting guidelines for DEI in your visualizations, we would love to hear about them.

REFERENCES

[1] Betzendahl, Lindsay. "Racial Inequality in the United States." Data visualization. Tableau Public, June 24, 2020. https://public.tableau.com/en-us/gallery/racial-inequality-united-states?tab=viz-of-the-day.

[2] Blakely, Pieta. "Presenting Data for a Targeted Universalist Approach." March 6, 2019. http://blog.pietablakely.com/presenting-data-for-a-targeted-universalist-approach.

[3] Borkin, Michelle A., Zoya Bylinskii, Nam Wook Kim, Constance May Bainbridge, Chelsea S. Yeh, Daniel Borkin, Hanspeter Pfister, and Aude Oliva. "Beyond Memorability: Visualization Recognition and Recall." *IEEE Transactions on Visualization and Computer Graphics* 22, no. 1 (August 2015): 519–28.

[4] Boy, Jeremy, Anshul Vikram Pandey, John Emerson, Margaret Satterthwaite, Oded Nov, and Enrico Bertini. "Showing People behind Data: Does Anthropomorphizing Visualizations Elicit More Empathy for Human Rights Data?" *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (2017): 5462–74.

[5] Bui, Kim. "Designing Data Visualisations with Empathy." Datajournalism.com, June 11, 2019.

https://datajournalism.com/read/longreads/designing-data-visualisations-with-empathy.

[6] Costanza-Chock, Sasha. "Design Justice: Towards an Intersectional Feminist Framework for Design Theory and Practice." *Proceedings of the Design Research Society* (2018).

[7] Groeger, Lena V. "A Big Article about Wee Things." ProPublica, September 25, 2014. https://www.propublica.org/nerds/a-big-article-about-wee-things.

[8] Harris, Jacob. "Connecting with the Dots." Source, January 15, 2015. https://source.opennews.org/articles/connecting-dots/.

[9] Kim, Erika. "Visible Women: Gender in Iconography." Noun Project, October 29, 2019. https://blog.thenounproject.com/visible-women-gender-in-iconography-c6518ec1ddb8.

[10] Klein, Scott. "Intro: The Design and Structure of a News Application." GitHub, updated March 4, 2013. https://github.com/propublica/guides/blob/master/design-structure.md.

[11] PolicyViz. "Style Guides." November 30, 2016. https://policyviz.com/2016/11/30/style-guides/.

[12] Schwabish, Jonathan. *Better Data Visualizations: A Guide for Scholars, Researchers, and Wonks*. (New York: Columbia University Press, forthcoming).

[13] Shirley, Rachel. "How Colors Advance and Recede in Art." *Science of Color* (blog). Accessed July 23, 2020. https://sites.google.com/site/scienceofcolour/how-colors-advance-and-recede-in-art.

[14] Steenhout, Nicholas. "Disability Language Is a Nuanced Thing," *Part of a Whole* (blog). July 9, 2020. https://incl.ca/disability-language-is-a-nuanced-thing/.

[15] Urban Institute. "Tracking COVID-19's Effects by Race and Ethnicity." Updated July 23, 2020. https://www.urban.org/features/tracking-covid-19s-effects-race-and-ethnicity.

[16] Urban Institute. "Urban Institute Data Visualization Style Guide." Accessed July 23, 2020. http://urbaninstitute.github.io/graphics-styleguide/.

[17] Wartell, Sarah Rosen. "Turning Insights into Action to Increase Upward Mobility." *Urban Wire*, (blog). May 24, 2018. https://www.urban.org/urban-wire/turning-insights-action-increase-upward-mobility.

[18] WebAIM. "Contrast Checker." Accessed July 23, 2020. https://webaim.org/resources/contrastchecker/.