EDITOR: Theresa-Marie Rhyne, theresamarierhyne@gmail.com

DEPARTMENT: VISUALIZATION VIEWPOINTS

Rainbow Colormaps Are Not All Bad

Colin Ware, *University of New Hampshire, Durham, NH, 03824, USA*

Maureen Stone, *Tableau Research, Seattle, WA, 98103, USA*

Danielle Albers Szafir, *University of North Carolina at Chapel Hill, Chapel Hill, NC, 27599, USA*

Some 15 years ago, Visualization Viewpoints published an influential article titled Rainbow Color Map (Still) Considered Harmful *(Borland and Taylor, 2007). The paper argued that the "rainbow colormap's characteristics of confusing the viewer, obscuring the data and actively misleading interpretation make it a poor choice for visualization." Subsequent articles often repeat and extend these arguments, so much so that avoiding rainbow colormaps, along with their derivatives, has become dogma in the visualization community. Despite this loud and persistent recommendation, scientists continue to use rainbow colormaps. Have we failed to communicate our message, or do rainbow colormaps offer advantages that have not been fully appreciated? We argue that rainbow colormaps have properties that are underappreciated by existing design conventions. We explore key critiques of the rainbow in the context of recent research to understand where and how rainbows might be misunderstood. Choosing a colormap is a complex task, and rainbow colormaps can be useful for selected applications.*

If you open a newspaper or turn on the local news, you will likely see a weather map using some variant of a *rainbow colormap*—a set of continuous colors traversing the colors of the rainbow—to represent temperature (e.g., Figure 1). However, decades of design guidance in visualization laments the ineffectiveness of the rainbow for communicating continuous data like temperatures. Despite this guidance, while the use of rainbows has declined in the visualization research literature, it has not declined in domain science publications.[2] Users continuously demand tools that give them the option to apply rainbow colormaps to their data. What about rainbows makes them so appealing? Do rainbows better

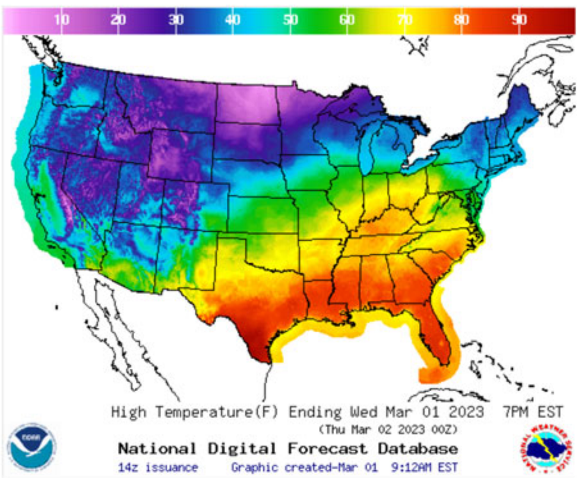

FIGURE 1. A temperature map of the United States from 1 March 2023. Produced by the National Weather Service of the National Oceanic and Atmospheric Administration (https://graphical.weather.gov/).

Digital Object Identifier 10.1109/MCG.2023.3246111
Date of current version 18 May 2023.

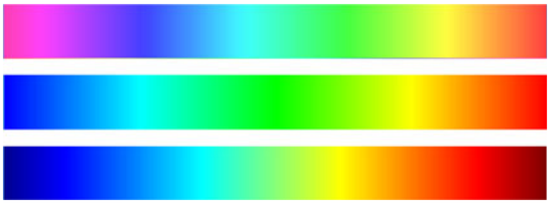

FIGURE 2. Three popular rainbow colormaps. Top: The classic HSV Rainbow. Middle: The Paraview Blue-Red Rainbow has no magenta. Bottom: Jet adds dark ends.

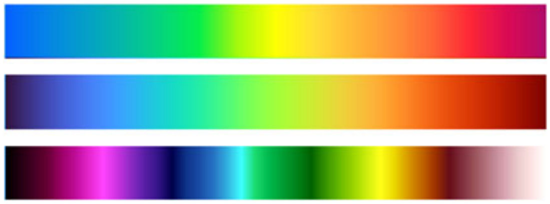

FIGURE 3. Top: A perceptually uniform rainbow with a diverging luminance profile (Paraview's Uniform). Middle: Turbo, which has a smoothly symmetric luminance profile. Bottom: A rainbow crafted for thermal imaging.[6]

support certain tasks, data, or scenarios that visualization guidelines fail to account for?

Visualization design heuristics tell us to avoid the rainbow. In fact, each of the authors of this article have written past articles discouraging the use of rainbow colormaps. There are three basic arguments against rainbow colormaps: 1) They are not monotonic in luminance, which can distort the shapes people infer from the data; 2) they are not perceptually uniform, which can distort the perception of data variation; and 3) they partition the data into distinctly colored regions or bands, which can create false divisions within the data. However, experimental evidence on the rainbow's utility is mixed. Rainbow colormaps have been repeatedly shown to be the most accurate in studies where a value must be read using a key.[3] Being able to read a temperature value is essential for common visualizations like weather maps, where people need to get an accurate estimate of the predicted temperature at specific locations. Rainbow colormaps may draw attention to global structure within the data[4] and enable people to better reason about relationships between data distributions.[5]

The critiques levied against rainbows are valid, especially for the original "rainbow": a simple interpolation of the display primary colors in many early visualization systems. Furthermore, full rainbow maps are very difficult to make accessible, especially to those with severe colorblindness. However, recent evidence calls on us to re-evaluate our dogmatic opposition to rainbows. We argue that many of the problems with rainbows can be mitigated through better design, and that for some tasks, these characteristics can actually be advantages, especially when rainbows produce beneficial emergent features that align well with the perceptual segmentation introduced by hue.

RAINBOW COLORMAPS

At its most basic, a colormap is a sequence of colors used to map data values onto an image. No colormap is an exact reproduction of the colors of an atmospheric rainbow that we encounter in nature. The term "rainbow" instead refers to an ordering of colors that matches that of the visible light spectrum—violet-blue-cyan-green-yellow-orange-red. While the classic rainbow traverses these colors using a transformation of the red, green, blue (RGB) display primaries to define hue, saturation, and value (HSV), many variants have been developed that vary in luminance, saturation, and hue to leverage different perceptual affordances that are useful for various visualization tasks.

Figure 2 shows the original HSV Rainbow and two more recent variants—the Paraview Rainbow and Jet—that drop the magenta endpoint of the original rainbow. Because they are simple interpolations in RGB (a perceptually nonlinear colorspace), these rainbows inevitably perform poorly on any test of perceived metric distance between displayed quantities. Despite these shortcomings, people often prefer these colormaps to best practice nonrainbow colormaps, either because they are familiar or because they provide value that we have yet to clearly identify. Rather than continuing to reiterate their known faults, research should aim to characterize when they seem to be effective as well as why people simply like them.

Many recent colormaps have been designed using perceptual principles to address known limitations in classic rainbows, namely to control luminance and provide spatial uniformity (Figure 3). Paraview's Uniform has a bivariate luminance profile like a diverging colormap. Turbo's profile is similar but smoothly rounded to minimize banding. The thermal rainbow intentionally leverages the banding effects caused by hue variations within the rainbow to highlight relevant regions in imaging data while manipulating local luminance to help people resolve smaller features.[6] The range of alternative rainbow colormap designs suggest that carefully crafted rainbows may not uniformly suffer from the same limitations that led to the visualization research community's universal antirainbow dogma. These designs may also

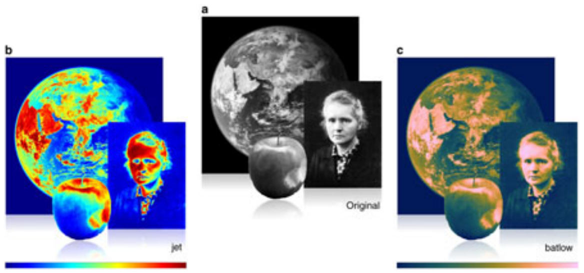

FIGURE 4. Three images demonstrating the effect of different colormaps from Crameri et al.[7]

explain the mixed performance of rainbows in recent empirical studies. In the ensuing sections, we discuss how thoughtful rainbow colormap design addressing three primary critiques of the rainbow—nonmonotonic luminance variation, perceptual nonuniformity, and hue banding—may reduce these limitations in practice to make rainbows more useful.

NONMONOTONIC LUMINANCE

Studies in vision science show that the shape of a surface can be inferred from luminance variation along the surface (i.e., shape-from-shading). However, we do not see shape-from-shading when there is only hue change in an image. For example, when the black and white portrait shown in Figure 4 is displayed using a rainbow, it seems very unnatural and confusing: we can see the face but much less easily than in the two photographs on the right. This difficulty is not necessarily because rainbows use hue, but because the rainbow varies irregularly in luminance. The Batlow colormap on the right varies in hue but is also monotonic in luminance, so the photograph appears relatively natural.

This and similar tests have been used to argue against rainbow colormaps,[7,8] which typically lack monotonic luminance, but they may not be an entirely fair test for visualization. The curved surfaces present in the apple and portrait rely heavily on shape-from-shading to resolve their structure. Our visual system expects changes along these surfaces to be reflected in the image's luminance values. If we were to do the reverse and take the hue values of a photograph and convert them to luminance, our visual system may run into similar challenges.

Removing the association between luminance and curvature impedes our ability to perceive shape-from-shading information in ways that may not inhibit finding patterns in abstract imagery like visualizations to the same degree. For example, in Figure 4, the apple and the portrait look strange when mapped to a rainbow. However, the temperature map of the Earth does not. Interpreting a temperature distribution does not rely on perceived shape-from-shading information. Perceived curvature from the grayscale and Batlow colormaps are not necessary to preserve in the maps. People can correctly infer the color semantics of the rainbow to intuitively make judgments about temperature. Losing the ability to perceive shape-from-shading in rainbows is not relevant when shape-from-shading perception is not required. Instead, what is most important is whether we can properly perceive and interpret critical patterns in the data. For example, rainbows outperform grayscale for comparing data distributions[5] and other tasks involving more global properties of data.[4] These tasks may rely on a different set of perceptual mechanisms than shape-from-shading. Future research into the mechanisms underlying our abilities to leverage rainbows in these contexts will help designers understand when and where rainbows become useful.

Many information visualizations that do not require people to resolve structural details in data, such as scatterplots or line graphs, use colormaps with monotonic luminance as people naturally associate "darker" with "more" on a white background, providing an intuitive mapping of relative quantities.[9] Sequential colormaps are therefore intentionally monotonic in luminance, and diverging maps like cool–warm are light in the center and dark on each end. Multihued monotonic colormaps such as Viridis or Magma may leverage similar mechanisms to rainbows while providing monotonic colormaps, but creating fully monotonic rainbows is difficult because both blue and red are dark colors on a display. However, options like Turbo or Paraview's Uniform offer rainbow colormaps with diverging luminance profiles that may preserve the benefits of both the hue variations of rainbows and the luminance variations of diverging colormaps.

PERCEPTUAL NONUNIFORMITY

Color science describes the smallest difference between two colors that people can detect as a *just noticeable difference* (JND). The perceptual uniformity of a colormap refers to how many JNDs there are between adjacent colors in the sequence. If a colormap is uniform, there will be an equal number of JNDs between equal steps in the color sequence. The number of JNDs at a particular range on a colormap also tells us how well people can resolve features using the colors in that range: if there are more JNDs between successive colors, people will be able to resolve more features in their data using that section of colormap.

Perceptual uniformity seems straightforward, but the number of JNDs between a pair of colors depends on many perceptual factors, including the colors

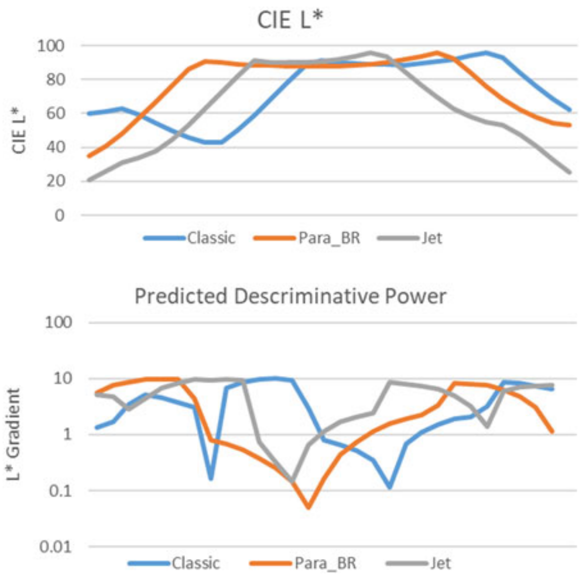

FIGURE 5. Top: The scaled luminance profiles of the rainbow colormaps in Figure 2. Bottom: The gradient of the scaled luminance profiles predicts resolving power for small features.

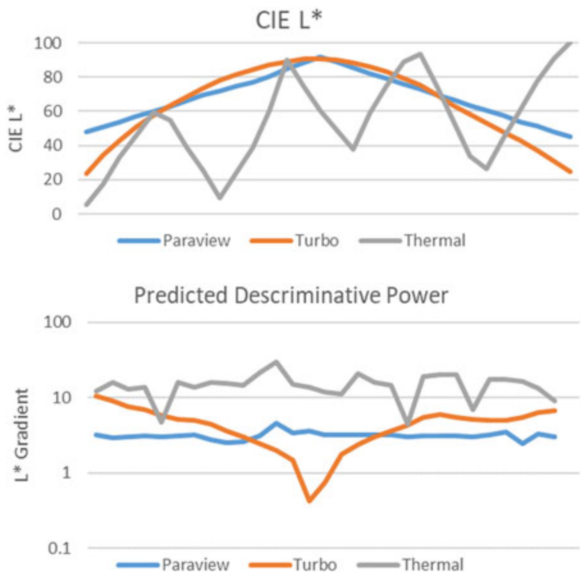

FIGURE 6. Top: The scaled luminance profiles of the rainbow colormaps in Figure 3. Bottom: The gradient of the scaled luminance profiles predicts resolving power for small features.

surrounding the region of interest, the size of the colored mark or region, and the state of adaptation of the viewer. JNDs are less useful measures at larger color distances (e.g., between green and red) and are sensitive to spatial variation (e.g., different patterns and textures that may emerge in continuous data). The smaller the feature in the data, the less color (in the sense of hue and saturation) contributes to our abilities to detect those features relative to luminance.[6] This correlation means that for a colormap to effectively show details in the data, it should have a lot of luminance variation. In other words, feature resolving power depends mostly, although not entirely, on luminance variation.

We can apply these ideas to understand tradeoffs in different rainbow colormap designs. CIE L* is a measure of perceived luminance often used to create and evaluate colormaps. The L* profiles of the colormaps shown in Figures 2 and 3 are shown in Figures 5 and 6. The gradient of L* at a particular point on a colormap is highly correlated with how well we can detect features in visualized data (hue and saturation also contribute to a limited degree).

We can use L* to show how the feature resolving power of a colormap varies along its length, as shown in Figures 5 and 6. Figure 5 shows that classic HSV Rainbows are extremely nonuniform in terms of how much data they communicate at different points, especially between cyan and green. Rainbows can resolve this issue by creating a diverging rainbow with a luminance peak in the center and linear (in L*) darkening to the left and right, as shown at the top in Figure 3. This rainbow has relatively uniform feature resolving characteristics. Flat spots in the luminance profile, as in the center colormap in Figure 3, lead to poor feature resolving power in this section.

Designers can also increase the feature resolving power of a colormap by introducing waves in the colormap's luminance profile. For example, the bottom colormap in Figure 3 traverses from light to dark (or the reverse) seven times. It has much greater feature resolving power than any monotonic colormap. However, the tradeoffs of this design choice are not yet well understood.

HUE BANDING

One of the most significant criticisms of the rainbow colormap is that the different hues create a visual segmentation: a green region, a red region, a blue region, etc. This strong effect is useful if these categories are meaningful, such as mapping temperatures above and below freezing to two different hues. However, people tend to read meaning into color categories even when they are simply artifacts of the colormap. Critics of the rainbow colormap rightfully worry about this problem. Meaningful segmentation requires controlling the function that maps data to colors to intentionally align data semantics with perceived bins, possibly even as part of the data-to-image generation process. Most discussions of colormaps in scientific visualization do

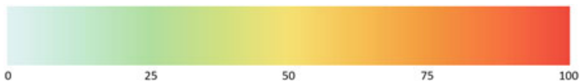

FIGURE 7. A rainbow designed for Tableau to visualize density. Hues fall on quartile boundaries: blue=0%, green=25%, yellow=50%, orange=75%, red=100%.

not address this step, and few applications offer convenient ways to do so.

Aligning data semantics with visual semantics has received more attention in information visualization. Categorical palettes use hue for rapid identification, ideally aligned with data semantics.[10] For quantitative data, a deliberate mapping between data and color breaks is common, especially in cartography.[11] Typically, these classified data are mapped to a stepped or banded colormap.

Statistical properties, such quantile boundaries, can be useful to highlight meaningful regions within a distribution. For example, Figure 7 shows a rainbow colormap used to enhance Tableau scatterplots by mapping data density to color. It is designed to emphasize areas of high density with warm, dark colors that standout against the cool, light blues and greens (e.g., the hotspots in Figure 8). This rainbow variant is visually monotonic, and primary hues are placed at quartile boundaries to act as a "color ruler." Its luminance profile is uniform and monotonic except for small variation around green, which was easier to identify when made slightly darker than the yellow.

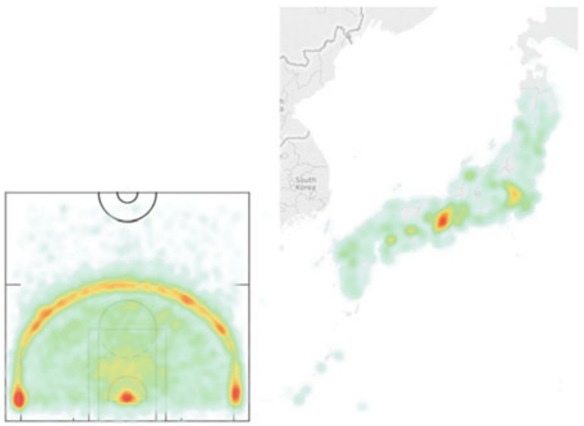

FIGURE 8. Tableau heatmap images using the rainbow in Figure 7. The first corresponds to basketball shots, plotted by position. The three point line and under the basket show as unsurprisingly dense. The second plots crime reports on a background map; high-crime areas are highlighted in red-orange (the large red spot is Tokyo).

Hue segmentation might also help manage cognitive load in complex visualization tasks by grouping subranges in the data, in effect creating implicit hierarchical structures people can reason across. Preliminary studies suggest that the nameability of color regions may support such reasoning processes.[4] We speculate that this grouping may simplify the data to allow people to more effectively complete tasks involving large subsets of the data, such as comparing distributions,[5] even in cases where the data lack a meaningful mapping to hue ranges. Further research is needed to fully understand how the visual system simultaneously processes hue and luminance variation in colormaps to support a range of tasks.

CONCLUSION

We hope we have persuaded you that rainbows are not all bad. While poorly designed rainbow colormaps can correspond to poor data visualizations, we believe well-designed rainbows, like well-designed multi-hued colormaps, can be valuable tools for visualization.

We suspect that the main reason why many scientists choose a rainbow colormap for their default is that it provides a reliable, familiar, and accurate measuring device. With a key, they can obtain the quantity associated with any particular feature and can easily describe regions in the data using nameable colors. In addition, it surfaces patterns having the same value, maxima, and minima, and, if the right version of the rainbow is chosen, can have excellent feature resolving power. Some studies show that even more traditional rainbows can be effective for the right task.[3,4]

Classic rainbows have poor perceptual properties with respect to perceptual uniformity and lack a linear, or even a predictably varying, luminance profile. Rainbows like those in Figure 3 have been designed to overcome these limitations. We would like to see these well-designed rainbows used more often in evaluative studies, especially for tasks where luminance and uniformity are known to be important. It is unclear what value we gain by continuing to rediscover that classic rainbows do poorly.

A recent study[12] demonstrates that well-designed multihue maps can increase accuracy and performance in data visualization tasks, compared to single hue sequential and diverging colormaps. A well-designed rainbow is simply a multihued map with more colors, which suggests it could perform similarly well. Furthermore, aligning color segments with data semantics, such as quantiles, should reinforce both the estimation of specific and relative data values. We believe that for all forms of data visualization, making it easier to construct a more deliberate data-to-color correspondence would make all hue-varying colormaps, including the rainbow, much more powerful.

Accessibility is a more difficult problem to solve for rainbows. While more severe forms of colorblindness effectively reduce color perception to two hues plus luminance, those with milder forms can distinguish many of the different rainbow hues, even red and green. But all forms leave the viewer with accurate luminance perception. Colormaps that only encode with luminance, such as the Tableau density rainbow, are accessible as long as the information provided by color segmentation is not essential. Similarly, the designers of Turbo have found that it is more accessible than the classic rainbow, crediting its smooth, bivariate luminance profile. We believe there is much more to learn about creating and evaluating accessible rainbows.

We still have much to understand about how people perceive and use rainbows in practice, especially in discovering the tasks and contexts where they excel. There may be many, including as-of-yet unsuspected, reasons why people prefer rainbows to more "correctly" designed colormaps. We understand well enough why rainbows can be bad; let us focus instead on finding out when and why they are good.

ACKNOWLEDGMENTS

This work was supported in part by the NSF under Grant IIS-2046725 and in part by NOAA under Grant NA20NOS4000196.

REFERENCES

1. D. Borland and R. M. Taylor II, "Rainbow color map (still) considered harmful," *IEEE Comput. Graphics Appl.*, vol. 27, no. 2, pp. 14–17, Mar./Apr. 2007.
2. I. Gołębiowska and A. Çöltekin, "What's wrong with the rainbow? an interdisciplinary review of empirical evidence for and against the rainbow color scheme in visualizations," *ISPRS J. Photogrammetry Remote Sens.*, vol. 194, pp. 195–208, 2022.
3. C. Ware, "Color sequences for univariate maps: Theory, experiments and principles," *IEEE Comput. Graphics Appl.*, vol. 8, no. 5, pp. 41–49, Sep. 1988.
4. K. Reda, A. A. Salvi, J. Gray, and M. E. Papka, "Color nameability predicts inference accuracy in spatial visualizations," in *Comput. Graphics Forum*, vol. 40, no. 3, pp. 49–60, 2021.
5. K. Reda and D. A. Szafir, "Rainbows revisited: Modeling effective colormap design for graphical inference," *IEEE Trans. Vis. Comput. Graphics*, vol. 27, no. 2, pp. 1032–1042, Feb. 2021.
6. C. Ware, T. L. Turton, R. Bujack, F. Samsel, P. Shrivastava, and D. H. Rogers, "Measuring and modeling the feature detection threshold functions of colormaps," *IEEE Trans. Vis. Comput. Graphics*, vol. 25, no. 9, pp. 2777–2790, Sep. 2019.
7. F. Crameri, G. E. Shephard, and P. J. Heron, "The misuse of colour in science communication," *Nature Commun.*, vol. 11, no. 1, pp. 1–10, 2020.
8. B. E. Rogowitz and A. D. Kalvin, "The which Blair project: A quick visual method for evaluating perceptual color maps," in *Proc. Vis.*, 2001, pp. 183–556.
9. K. B. Schloss, C. C. Gramazio, A. T. Silverman, M. L. Parker, and A. S. Wang, "Mapping color to meaning in colormap data visualizations," *IEEE Trans. Vis. Comput. Graphics*, vol. 25, no. 1, pp. 810–819, Jan. 2019.
10. V. Setlur and M. C. Stone, "A linguistic approach to categorical color assignment for data visualization," *IEEE Trans. Vis. Comput. Graphics*, vol. 22, no. 1, pp. 698–707, Jan. 2016.
11. C. Brewer and L. Pickle, "Evaluation of methods for classifying epidemiological data on choropleth maps in series," *Ann. Assoc. Amer. Geographers*, vol. 92, no. 4, pp. 662–681, 2002.
12. Y. Liu and J. Heer, "Somewhere over the rainbow: An empirical assessment of quantitative colormaps," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, 2018, pp. 1–12.

COLIN WARE is a professor emeritus with the University of New Hampshire, Durham, NH, 03824, USA. His research involves the application theories of perception to methods for data visualization. Ware received his Ph.D. in psychology from the University of Toronto. He is the corresponding author of this article. Contact him at cware@ccom.unh.edu.

MAUREEN STONE, retired, was with Tableau Research, Seattle, WA, 98103, USA. While best known for her expertise in digital color, she has a broad experience in information visualization, interactive graphics, and user interface design. Contact her at mstone@acm.org.

DANIELLE ALBERS SZAFIR is with the University of North Carolina, Chapel Hill, NC, 27599, USA. Her research explores the interplay of visualization and cognitive science. Szafir received her Ph.D. in computer sciences at the University of Wisconsin-Madison. She is a member of the IEEE. Contact her at danielle.szafir@cs.unc.edu.

Contact department editor Theresa-Marie Rhyne at theresamarierhyne@gmail.com.