



© 1995 IBM Corporation



How NOT to Lie with Visualization

Bernice E. Rogowitz

`rogowitz@watson.ibm.com`

Lloyd A. Treinish

`lloyd@watson.ibm.com`

**IBM Thomas J. Watson Research Center
Yorktown Heights, NY**

Introduction

How data are represented visually has a powerful effect on how the structure in those data is perceived. For example, in [Figure 1](#), four representations of an MRI scan of a human head are shown. The only difference between these images is the mapping of color to data values, yet, the four representations look very different. Furthermore, the inferences an analyst would draw from these representations would vary considerably. That is, variations in the method of representing the data can significantly influence the user's perception and interpretation of the data.

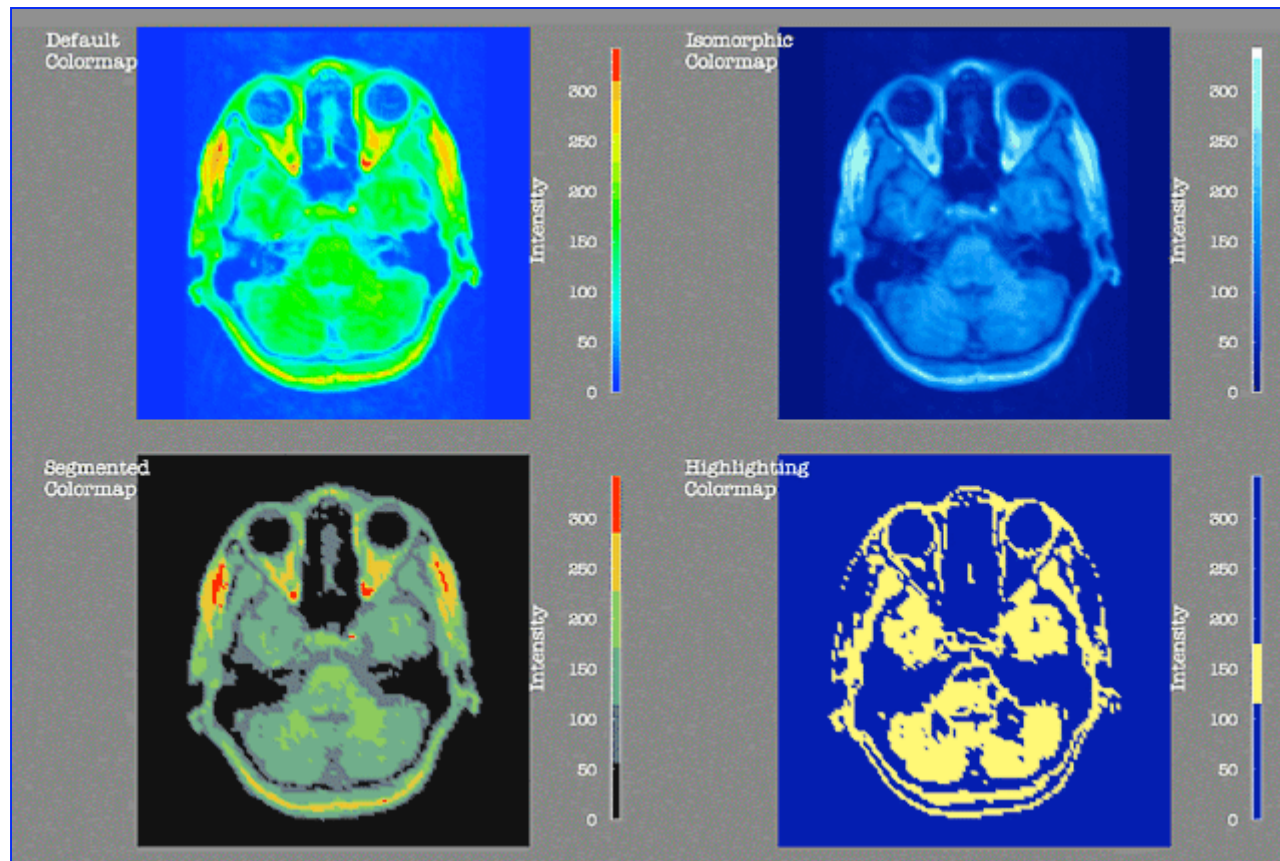


Figure 1. Four colormaps applied to a slice of an MRI scan of a human head. They demonstrate how the representation can influence the interpretation of the data.

The importance of visual representation has been a lively topic at the annual IEEE Computer Society Visualization conferences. This concept was first publicized by Huff [1954] in his book *How to Lie with Statistics*. In this book and in the *How to Lie with Visualization* sessions at those conferences, the major concern is how the interpretation of data can be subverted by manipulating the data representation. In this article, we take a converse tack, and ask -- how can the interpretation of data be enhanced? To address this question, we consider the structure of the data, the perception of the visual dimensions used in visualization, and the task the analyst is trying to solve.

The Complexity of Mapping Data on Visual Representations

Modern interactive systems give the user free reign over the mapping of data onto visual dimensions, and the number of visual dimensions available for data representation is exploding. A visualization can use x, y, and z to represent the spatial dimensions of an object, color can be mapped onto a surface representing a fourth, the surface can be deformed according to a fifth, isocontour lines can represent a sixth, coloring them can represent a seventh, glyphs on the surface can represent a few more, not to mention animation, transparency, and stereo. This great flexibility, however, can open a Pandora's box of problems for the user, and can easily give rise to visual representations which do not adequately represent the structure in the data or which introduce misleading visual artifacts.

The appropriate use of color is an area of particular consternation. This is partly because the perceptual impact of a color cannot be reliably predicted from a knowledge of the red, green and blue components generally made available to users. Furthermore, even if the three perceptual dimensions of color are surfaced to the users, they may not be aware that different aspects of the color signal communicate different characteristics of the data. Without guidance about the physical or psychophysical properties of color, or about which colormaps are most appropriate for which types of data, the user is at a loss, even if the system provides a colormap editor or a library of pre-computed colormaps.

One common way developers of visualization software address this problem is to provide users with a default colormap. The most common default colormap, shown in the top left panel of Figure 1, maps the lowest value in the variable to blue, the highest value to red, and interpolates in color space (red, green, blue) to produce a color scale. This rainbow hue colormap is widely used in visualization, but produces several well-documented artifacts (e.g., Lefkowitz and Herman [1992]; Robertson [1988]; Rogowitz, Ling and Kellogg [1992]). In this MRI image, for example, the colormap creates perceived contours which do not reflect discrete transitions in the data, structures in the data which fall within one of these artificial bands are not represented, and attention is drawn to the yellow areas because they are the brightest, not because they are in any way the most important.

Giving users tools for creating accurate and effective representations of their data, furthermore, is not confined to the selection of colormaps. There is also confusion in the application of contours, transparency, depth, and animation, especially since the perception of these dimensions can often interact. For example, if a blue and a red object are placed behind a translucent green object, you might expect that both objects would maintain their color, but be tinged by the color of the transparent layer. This is true for the blue object, which appears bluish-green, but not for the red object, which appears yellow. This effect is well understood within the context of the algorithm by which transparency is generally computed and principles of additive color mixture, but can produce surprising results for the user (Rogowitz and Treinish [1993a]).

Using Perceptual Rules to Guide the Design Process

Since most users do not want to become experts in human perception, our strategy is to incorporate guidance directly into the visualization software to aid in the the visual design process (Rogowitz and Treinish [1993a]; [1993b]; [1994]; and Bergman, Rogowitz and Treinish [1995]). In our approach, which we call PRAVDA (Perceptual Rule-based Architecture for Visualizing Data Accurately), rules filter the choices offered to the user, based on principles of human perception, attention, and color theory.

In the case of colormap selection, for example, we have constructed a library of colormaps and a set of perceptual rules which constrain the set of colormaps offered to the user. These rules are parameterized by metadata about 1) data type 2) data spatial frequency, 3) visualization task, and 4) other design choices made by the user. Three colormaps designed for different visualization tasks are compared with the default colormap in [Figure 1](#). The isomorphic colormap (upper right) is designed to produce a faithful representation of the structure in the data. In this isomorphic colormap equal steps in data value correspond to equal perceptual steps in the color scale. The segmented colormap (lower left) is designed to delineate regions visually. The highlighting colormap (lower right) is designed to draw the users' attention to regions in the image which have certain characteristic features (lower right). This color map was designed to draw attention to areas which have data values near the median of the range.

The four colormaps in [Figure 1](#) clearly demonstrate how different mappings of data onto color scales produce different representations of the data. The goal of our work is to understand how different information in the data is communicated by specific characteristics of the visual representation, and to harness this knowledge so it can be used routinely in visualization. The rest of this article focuses on the colormap problem, describing the perceptual rules and metadata required to drive colormap selection.

Faithfully Representing the Structure in the Data

In order to accurately represent the structure in the data, it is important to understand the relationship between data structure and visual representation. For nominal data, objects should be distinguishably different, but since the data themselves are not ordered, there should be no perceptual ordering in the representation. For ordinal data, objects should be perceptually discriminable, but the ordering of the objects should be apparent in the representation. In interval data, equal steps in data value should appear as steps of equal perceived magnitude in the representation. In ratio data, values increase and decrease monotonically about a true zero or other threshold, which should be preserved in the data representation.

One important application of scientific visualization is to represent the magnitude of a variable at every spatial position. In many cases, the interpretation of the data depends on having the visual picture accurately represent the structure in the data. In order to accurately represent interval data, for example, the visual dimension chosen should appear continuous to the user. Candidate colormaps which preserve the monotonic relationship between data values and perceived magnitude can be drawn from psychophysical scaling experiments. Stevens [1966], for example, identified a set of sensory dimensions for which a monotonic increase in stimulus intensity produced a monotonic increase in perceived magnitude. In particular, he found the shape of this relationship to be a power law, with each sensory dimension characterized by its exponent. Perceived magnitude obeys a

power relationship with physical luminance over a very large range of gray scales, which may explain why grayscale colormaps are commonly used in medical imaging. Another dimension which displays this behavior is color saturation, the progression of a color from vivid to pastel.

The top row of [Figure 1](#) compares the effectiveness of the default colormap and a colormap designed to produce an isomorphic representation of interval data. Looking at the color bar for the default colormap we see bands of colors, not a gradual increase across the range. For example, nearly the entire range from 50 to 100 looks uniformly cyan. Although the data change by almost a factor of 2, all the values in this range look identical. This is also true for magnetic resonances in the range from 125 to 200, which appear to be green. This colormap produces a contoured impression, and masks the subtle variations in MRI intensity.

In contrast, the isomorphic representation used in the upper right, although less dramatic, more accurately reflects the underlying structure in the data. In this colormap, luminance and saturation both increase monotonically with data value. That is, brightness increases monotonically and hue, which begins as a pure vivid blue, becomes more and more pastel. This colormap produces a monotonic increase in perceived magnitude over the range. Using this colormap, structures which are invisible using the rainbow hue map can be easily seen. For example, the spatial structure in the midbrain and striate cortex which appear uniform green in the default map are highly detailed in the isomorphic map. Given the artifacts introduced by the default colormap, it is easy to understand why the medical community has been so cautious about adding color to their visual representations.

The Importance of Spatial Frequency

Not all isomorphic colormaps, however, are appropriate for all data sets because different components of the color signal are processed differently by the human visual system. One very important distinction for visualization is that the components of the color signal have different spatial sensitivities. The luminance component in a color (the brightness/darkness component) is critical for carrying information about high spatial frequency variations in the data. If the colormap does not contain a monotonic luminance variation, fine resolution information will not be seen. Conversely, the saturation and hue components in color are critical for carrying information about low spatial frequency variations in the data. A colormap which only varies in luminance (e.g., a grayscale image) cannot adequately communicate information about gradual changes in the spatial structure of the data.

This means that the balance of luminance and saturation variation in an isomorphic colormap depends on the spatial frequency of the data. Interval data with high spatial frequency information call for a monotonic scale with a strong luminance component; interval data with low-spatial-frequency information call for a monotonic scale with a strong saturation component.

These ideas are illustrated in [Figure 2](#), which shows a luminance-based colormap (left side) and saturation-based colormap (right side) applied to low-spatial-frequency data (top) and high-spatial-frequency data (bottom). In all four cases, continuous data are mapped onto isomorphic colormaps, so contouring and other artifacts have already been eliminated. This figure thus highlights additional advantages of taking spatial frequency into

account.

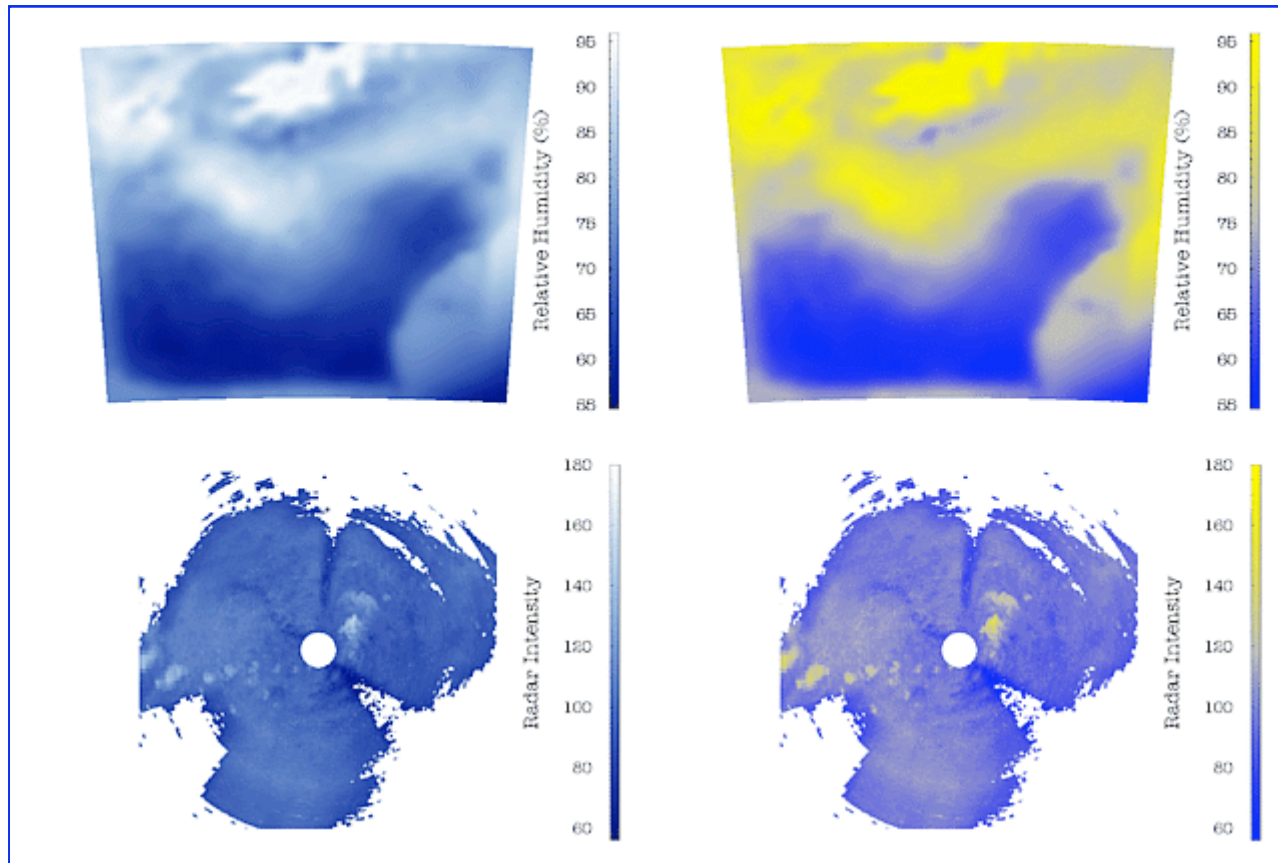


Figure 2. Isomorphic colormaps for low and high spatial frequency data. The top row shows low spatial frequency data from a weather model. The bottom row shows high spatial frequency data from a radar scan. The high frequency colormap (left) reveals more detail in the radar data. The low frequency colormap (right) reveals more structure in the weather data.

The data in the top row are output from a weather model which computes, among other things, the variation in relative humidity over a geographic region. The structure of this low spatial-frequency variation is practically lost when the data are depicted with a map designed for depicting high

spatial frequency information (top left). The right-hand map, designed to expose low spatial-frequency structure, gives the analyst more information, especially in regions where the humidity changes slowly over the geography such as near the lower central portion below 65%. Also, in the lower right-hand corner of the image, the infusion of the high humidity air into the low-humidity area is clearly seen as a yellow stream, virtually invisible with the colormap intended for high-spatial-frequency data.

The images in the bottom row show a radial sweep from a weather radar sensor, measuring the high-spatial frequency variation of reflected intensity (e.g., from thick clouds). The high spatial frequency map (left) gives a good representation of the finely detailed structure of these data, and also reveals sampling artifacts introduced by the sensor. The low spatial frequency colormap (right) blurs the fine detail and, because the values above the mean are a different hue, puts inappropriate emphasis on these regions, shown in yellow.

Colormaps for Segmentation Tasks

The rules for providing isomorphic colormaps for ratio and interval data are also effective in creating maps for segmenting data. The luminance component conveys monotonicity for high spatial frequency data, while the saturation component can be used to convey monotonicity in low spatial-frequency data. Since the steps are explicitly defined, however, luminance steps can also be effectively used for low spatial-frequency data. In creating a segmented colormap, it is important that the segments are each discriminably different from one another, which limits the number of steps which can be represented. We have found that more steps can be effectively discriminated for low spatial-frequency data than for high.

[Figure 3](#) shows a five-level segmented colormap (left-side) and a ten-level segmented colormap (right-side) applied to low-spatial frequency data (top) and high-spatial frequency data (bottom). For low spatial-frequency data (top row), additional levels provide additional information. In this case, additional features of the earth's magnetic field in the southern hemisphere are revealed. For example, in the right-hand image, the gradient about the south magnetic pole is clearer. By contrast, additional features of the high spatial-frequency cloud fraction observations (bottom row) are not revealed by increasing the number of colormap steps, effectively blurring the segmentation.




Figure 3. Segmented colormaps for low and high spatial frequency data. The top row shows low spatial frequency data of the earth's magnetic field. The bottom row shows high spatial frequency cloud fraction data. The high frequency colormap (right) reveals more information about the structure of the low-frequency data, but reduces the information communicated for the high-frequency data.

Colormaps for Highlighting Tasks

Rules for selecting colormaps which highlight particular features in the data can be drawn from the literature on attention (e.g., Treisman and Gelade [1980]; Julesz [1981]). For visualization, this requires a user to identify ranges of data to highlight perceptually such as shown in the lower right of

[Figure 1.](#)

Using these principles, it is possible to construct colormaps which highlight particular ranges in the data. An interesting extension of this approach is illustrated in [Figure 4](#), which displays data from the visible part of the spectrum remotely-sensed from space.




Figure 4. Two isomorphic colormaps applied to remotely sensed data, and a highlighting scheme which identifies regions of interest without disturbing the perceived spatial structure of the data.

The left-hand panels display these data using two isomorphic colormaps designed for high-spatial-frequency data. The right-hand panel shows how color can be used to highlight a region of interest without disturbing the perception of other aspects of the data. Across the entire image, the luminance component of the colormap is identical. Within the regions of interest, however, the hue component is varied, producing three distinct, semantically differentiable regions, one blue, one green and one yellow. This method has been used successfully to help users mark regions of interest and to highlight for the user regions in the image which display certain characteristics, such as containing data which match a template.

Complementary Visual Techniques

An important task in visualization is to represent data from many sources simultaneously. The image at the top of Figure 5 is derived from three

spectral bands of another remotely-sensed image. These data are displayed in a typical fashion, mapping the values of each spectral band to levels of red, green and blue in the image. This representation provides a crude classification of the pixels. Each pixel in these data have also been categorized into 5 classes using an external land-use classification scheme. This information could be displayed to the user by coloring the pixels according to class membership, with a different color for each class.




Figure 5. Two approaches to studying the results of a land-use classification model for remotely sensed data. The top figure shows a typical

pixel-based colormap. The bottom figure illustrates a graphical approach to examining differences between classes with respect to two spectral bands.

If the spatial regions occupied by the classes are sufficiently large, each pixel could be mapped onto isoluminant blue, cyan, green, yellow, and red, as described above, to highlight the different categories without perturbing the spatial structure of the data.

The six panels on the bottom illustrate a complementary method for using color to understand the semantics of class membership for such a data set. In this representation, each pixel has also been assigned a color according to its class membership. The coloring is used, however, to study the behavior of the different classes in terms of relationships among the various spectral bands. The top left plot shows the relationship between IR1 and IR2, the near-infrared and far-infrared bands. These bands are highly correlated ($r=0.92$). The next five plots show this same relationship separately for each of the five classes. Even looking at this one bivariate relationship reveals that the red and blue classes are different from the whole population and from the other classes in that there is a much smaller correlation between these two bands, that the green and yellow classes are the only classes with low values in both infrared bands, and the cyan class is the only one with high values in both infrared bands. This type of analysis allows the user to gain insight into the semantics of class membership.

PRAVDAColor

[Figure 6](#) shows the PRAVDA rule-based colormap selection tool incorporated into an IBM Visualization Data Explorer (Abram and Treinish [1995]) program. In this visual program, data are imported into Data Explorer and flow into a module called PRAVDAColor. This tool computes metadata about the spatial frequency of the data and about data type (e.g., ordinal, interval or ratio), and asks the user to select the goal of the visual representation (e.g., isomorphic, segmentation, or highlighting) via a control panel widget.

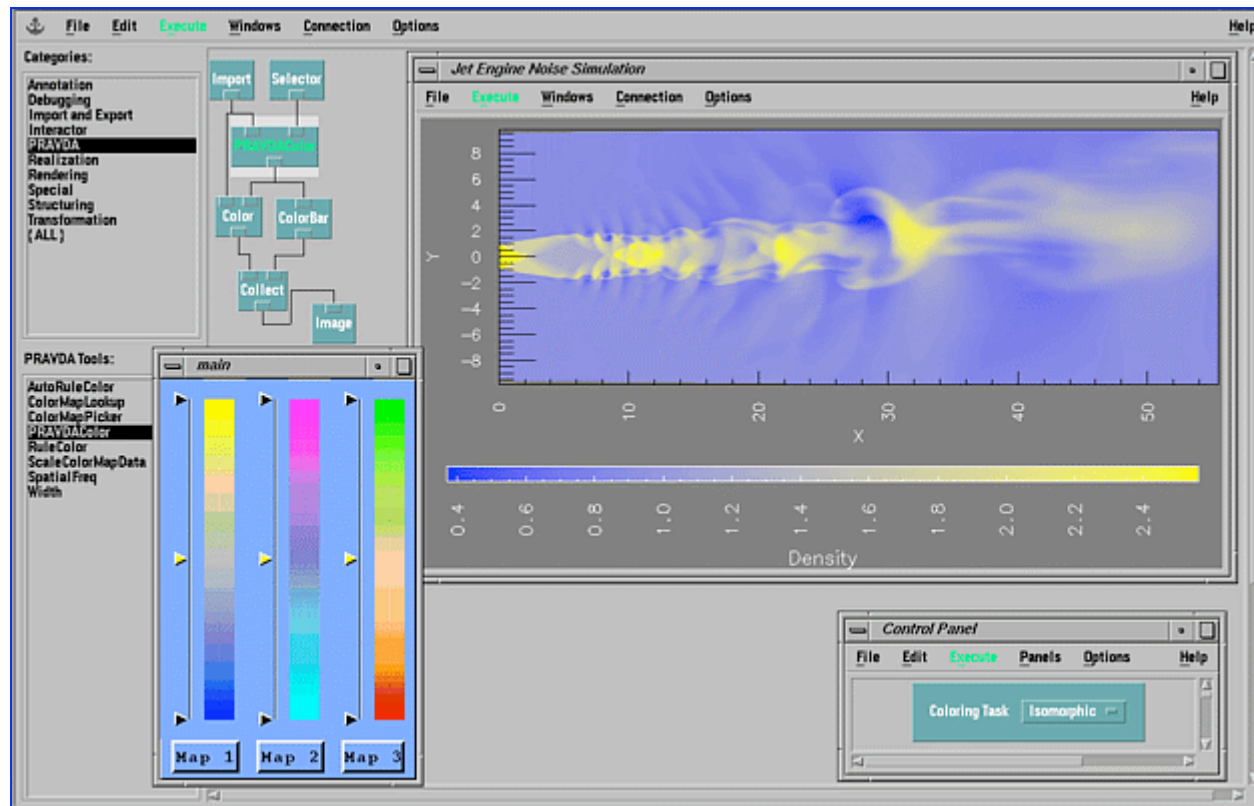


Figure 6. Data Explorer visual program incorporating PRAVDAColor, demonstrating selection of candidate isomorphic colormaps for a low-spatial-frequency data set.

These metadata flow to rules which constrain the set of colormaps offered to the user. In this case, three colormaps have been offered to the user. Since the simulated jet engine noise data shown in this example are low spatial frequency interval data, and the task selected is isomorphic, these colormaps all encode variations in magnitude as variations in the saturation of opponent-process pairs. Clicking on any of the colormaps applies them directly to the data, and the user is free to vary the range of the colormap. In this case, the full range of the first colormap has been selected and the data are represented by a blue/yellow saturation scale.

Conclusions

Modern systems for creating visualizations have evolved to the extent that non-experts can create meaningful representations of their data. However, it is still not easy enough, mainly because the visual effects of processing, realizing and rendering data are not well-understood by the user, and the process of creating visualizations is largely ad hoc. Often countless iterations are spent trying to get a color right, to draw attention to a particular juxtaposition in the data, or to understand why a feature on the display screen does not seem to correlate to a physical phenomenon.

Our approach emphasizes a migration from a tool-based visualization system to a rule-based system which helps the user navigate through a complex design space. Since the design process is iterative, the application of the rules is under interactive user control. The rules we have implemented so far draw on knowledge from the areas of human perception and color theory, but this structure could easily be extended to incorporate expertise from other domains. The goal of this system is to help users make better, faster representations of their data.

Acknowledgments

This work is partially supported under NASA grant CAN NCC5-101.

The authors wish to acknowledge Lawrence Bergman's contribution to the development of PRAVDAColor and John Gerth's extension of the PRAVDAColor isomorphic maps to region-of-interest highlighting in satellite images. We would also like to thank Vittorio Castelli and Ed Kalin for assistance with the image classification work described in Figure 5.

MRI data are available courtesy of New York University, New York, NY. Cloud fraction and magnetic field data are available courtesy of NASA/Goddard Space Flight Center, Greenbelt, MD. Humidity data are available courtesy of NOAA Forecast Systems Laboratory, Boulder, CO. Radar data are available courtesy of Sigmat, Inc., Westford, MA. Jet engine noise data are available courtesy of CRAFT, Inc., Dublin, PA. Remotely sensed images are available courtesy of EDC, USGS, Sioux Falls, SD.

The analysis and visualization at the bottom of Figure 5 were produced using the Diamond software, which was developed at the IBM T. J. Watson Research Center and is marketed by [SPSS, Inc.](#)

All other figures were produced with the [IBM Visualization Data Explorer](#), which was developed at the IBM T. J. Watson Research Center.

References

Abram, G. and L. Treinish. An Extended Data-Flow Architecture for Data Analysis and Visualization. Proceedings of the IEEE Visualization 1995 Conference, October 1995, pp. 263-270.

Bergman, L. D., B. E. Rogowitz and L. A. Treinish. A Rule-based Tool for Assisting Colormap Selections. Proceedings of the IEEE Visualization 1995 Conference, pp. 118-125, October 1995.

Huff, D. How to Lie with Statistics, New York. Norton, 1954.

Julesz, B. Textons, the elements of texture perception, and their interactions. Nature, 290, 12, 91-97, 1981.

Lefkowitz, H. and G. T. Herman. Color Scales for Image Data. IEEE Computer Graphics and Applications, 12, n. 1, pp. 72-80, January 1992.

Robertson, P. K. Visualizing Color Gamuts: A User Interface for the Effective Use of Perceptual Color Spaces in Data Displays. IEEE Computer Graphics and Applications, 8, pp. 50-63, September 1988.

Rogowitz, B. E., D. T. Ling and W. A. Kellogg. Task Dependence, Veridicality, and Pre-Attentive Vision: Taking Advantage of Perceptually-Rich Computer Environments. Proceedings of the SPIE Symposium, 1666, Human Vision, Visual Processing and Digital Display III, pp. 504-513, February 1992.

Rogowitz, B. E. and L. A. Treinish. Data Structures and Perceptual Structures. Proceedings of the SPIE Symposium, 1913, Human Vision, Visual Processing and Digital Display IV, pp. 600-612, February 1993.

Rogowitz, B. E. and L. A. Treinish. An Architecture for Perceptual Rule-Based Visualization. Proceedings of the IEEE Visualization 1993 Conference, pp. 236-243, October 1993.

Rogowitz, B. E. and L. A. Treinish. "Using Perceptual Rules in Interactive Visualization". Proceedings of the SPIE Symposium, 2179, Human Vision, Visual Processing and Digital Display V, pp. 287-295, February 1994.

Stevens, S. S. Matching Functions Between Loudness and Ten Other Continua. Perception and Psychophysics, 1, pp. 5-8, 1966.

Treisman, A. and G. Gelade. A Feature Integration Theory of Attention. Cognitive Psychology, 18, pp. 643-662, 1980.

