Published in final edited form as:

IEEE Comput Graph Appl. 2014 ; 34(6): 83–87. doi:10.1109/MCG.2014.116.

Random Per-Element Luminance Modulation for Improved Visual Tracking

Russell M. Taylor II and Jonathan Harter

University of North Carolina at Chapel Hill

Random per-element luminance modulation is a cross-cutting technique that, when applied to a wide range of visualizations, improves the visual salience of details and tracking of individual elements. We and other researchers have successfully applied luminance modulation in specialized conditions to deal with overplotting in parallel coordinates and to produce dense flow textures. (For a look at other ways to deal with overplotting, see the sidebar.)

Other Ways to Deal with Overplotting

One way to deal with overplotting in scatterplots and parallel coordinates is to include only a random subset of the data elements in a specified region of interest.^{1,2}

Another way is to reduce the data density overall by displaying only a subset of the elements across the entire image. Enrico Bertini and Giuseppe Santucci studied this approach's effectiveness, comparing uniform sampling with a perceptual nonuniform sampling in a user study.³ This could be effectively combined with random luminance modulation (see the main article) to both reduce clutter and trace individual elements.

A third way is to use blurring to produce density estimates of the line distributions, rather than showing individual lines. David Feng and his colleagues combined this with the display of uncertainty, producing a subset of the curves in parallel coordinates that were certain, along with density estimates of all the curves.⁴

Malte Zöckler and colleagues used transparency and illumination to effectively display dense sets of streamlines.⁵ This approach worked well in their target cases of electric and magnetic fields with noncrossing lines.

References

1. Ellis G. The Sampling Lens: Making Sense of Saturated Visualisations. CHI '05 Extended Abstracts on Human Factors in Computing Systems (CHI EA 05). 2005:1351–1354.
2. Ellis G, Dix A. The Plot, the Clutter, the Sampling and Its Lens: Occlusion Measures for Automatic Clutter Reduction. Proc 2006 Working Conf Advanced Visual Interfaces (AVI 06). 2006:266–269.
3. Bertini E, Santucci G. Give Chance a Chance: Modeling Density to Enhance Scatter Plot Quality through Random Data Sampling. Information Visualization. 2006; 5(2):95–100.
4. Feng D, et al. Matching Visual Saliency to Confidence in Plots of Uncertain Data. Trans Visualization and Computer Graphics. 2010; 16(6):980–989. [PubMed: 20975135]
5. Zöckler M, et al. Interactive Visualization of 3D-Vector Fields Using Illuminated Streamlines. Proc 7th IEEE Conf Visualization (VIS 96). 1996:107–113.

Here, we describe the use of random per-element luminance modulation in five contexts. These examples demonstrate that this technique has a broad range of applicability, improving detail and shape perception in 2D, 3D, and *ND* visualizations of scalar, vector, and tensor fields.

ND Scalar Display Using Parallel Coordinates

Figure 1 shows a hydrodynamic simulation of high-energy particle collisions. With standard parallel coordinates, overplotting causes lines to merge into blocks of solid color (see Figure 1a). In contrast, luminance-modulated parallel coordinates improve the patterns' visual salience (see Figure 1b).¹ In particular, the line bundles shaped like an S, a U, and an inverted U in the middle of the third parameter (U_{wind}) are more salient. Tracing an individual line across multiple axes is also easier.

In their basic form, luminance-modulated parallel coordinates multiply the luminance of each parallel-coordinates polyline by a factor determined by a uniformly distributed random variable between 0.5 and 1. Each line remains fully opaque. Until overplotting occurs, luminance-modulated parallel coordinates look quite similar to regular parallel coordinates because the intensity variation isn't high enough to hide individual lines against the dark background. In an overplotted region, luminance-modulated parallel coordinates maintain the distinction of each polyline from its neighbors (whose luminance values are randomly modulated), preserving the capability of each polyline to indicate its orientation. This leaves hue free for its traditional use in brushing clusters within parallel coordinates.

Vector Display Using Texture

It has long been known that random luminance modulation can be part of an effective technique for constructing textures that reveal 2D flow magnitude. Line-integral convolution (LIC) blurs an underlying random-intensity texture,² and spot noise overplots random-intensity oriented line glyphs.³ Both approaches achieve high-density display of flow magnitude and orientation.

Timothy Urness and his colleagues added data-dependent saturation scales to random luminance, using hue to discriminate between multiple data-sets.⁴ This augmented the random-luminance vector display with data-driven scalar information. Urness and his colleagues extended this to display data from two vector fields.⁵

Victoria Interrante and Chester Grosch extended the use of random colors in LIC to 3D.⁶ They demonstrated the use of random hue variation in addition to halos to enhance the visual separation of neighboring streamlines. They also recommended using random luminance modulation when hue is used for other purposes.

As with parallel coordinates, the addition of per-element random luminance modulation enables 2D and 3D LIC to maintain the ability to see individual elements, even with densities that produce complete overplotting.

3D Vector Display Using Glyphs

In an early GPU-based particle-simulation system,⁷ we used random per-sphere saturation variation to make it easier to distinguish individual particles in the overall flow. Because the spheres were blue and blue doesn't contribute to luminance, this had the side effect of introducing luminance modulation.

To confirm luminance modulation's effectiveness, we produced an animation of cones following wind flow in a weather simulation dataset from Michael Kiefer in Sharon Zhong's group at Michigan State University. Figure 2 shows the first frame of animations without and with random per-cone luminance modulation. Adding the luminance modulation greatly improved the viewer's ability to track a cone as it moved between and among neighbors. It also made it clear which cone passed in front during an overlap, disambiguating relative depth.

3D Vector Display Using Streamlines

Figure 3 shows streamline-based visualizations of the previous weather dataset without and with random luminance modulation.

As with parallel coordinates, per-element luminance modulation of 3D streamlines increases each line's salience as that line passes through regions with many other lines. It also highlights the relative depth of crossing streamlines because the contrasting brightness of the one in front breaks the constant brightness of the one behind.

3D Tensor Visualization Using Hyperstreamlines

Figure 4 shows hyperstreamlines of the stress tensor from a quantum-chromodynamics simulation by Hannah Petersen in Steffen Bass's group at Duke University. The lines are colored by an isoluminant blue-gray-red map to indicate local gluon density. Figure 4b combines random luminance intensity with color variation to show additional data.

Discussion

The perceived brightness, which is related to luminance, is the color channel the human visual system uses to understand object shapes and detect small features.⁸ This implies that you must take care when adjusting it to avoid introducing artifacts in the perception of the objects' shapes. (This is why isoluminant colormaps are usually better for mapping data onto 3D surfaces.) It also means that the perceived brightness is the channel that provides the best separation between detailed objects in the scene, which is why we use it.

Per-element luminance modulation applies the same modulation to all of each element (glyph, streamline, or hyperstreamline). So, it avoids the negative impact on surface shape perception of using a luminance-varying color map on top of 3D geometry. When such maps cause luminance gradients in an element, they distort shape-from-shading information about the element.⁸ Our approach paints each element a constant color with slightly different shades between elements.

As we mentioned before, when overplotting occurs, random luminance modulation enables the viewer to perceive individual polylines. The same situation occurs for streamlines, hyperstreamlines, and other extended structures. As these structures begin to overlap, luminance gradients excite local-orientation sensors in the retina at the scale of the structures. This enables the sensors to maintain the ability to determine local orientation and track individual elements.

For small moving elements such as the cones in Figure 2b, luminance modulation lets viewers pre-attentively distinguish targets on the basis of their brightness. This reduces the number of potentially confusing targets in each region.

Modulation

For the examples in this article, we modulated each element's luminance from 50 to 100 percent by multiplying each (R, G, B) component by the same constant factor. This range was enough to show clear visual differences among elements while maintaining sufficient with in-element contrast to display shape and color information.

Limitations

Random luminance modulation outperforms data-driven luminance modulation in some contexts, but not all. We've showed that adjusting line luminance on the basis of one of the scalar variables in parallel coordinates indicates trends between that variable and others (see Figure 5).¹ This technique might also effectively indicate groupings in other visualizations. However, when nearby elements all have a similar luminance, distinguishing neighbors again becomes challenging.

When applied to objects with color maps displaying secondary variables, luminance modulation shifts the color values at different locations by different amounts. To the extent that varying brightness distorts the perception of hue or saturation, this will cause perceptual shifts in the values read from these maps. We restricted such distortion in the visualizations in this article by limiting the modulation to keep the brightness at least 50 percent of the original.

Random luminance modulation has the unintended side effect of making some elements in the image brighter than the others and thus more salient, perhaps drawing undue attention to them. Fortunately, this is done in unbiased manner owing to the random selection.

Because of our success with random per-element luminance modulation, we continue to explore new cases in which it could improve visualization. We believe that it's a general technique that all visualization designers should consider.

Acknowledgments

We thank our collaborators Michael Kiefer, Sharon Zhong, Hannah Petersen, and Steffen Bass for the use of their datasets as examples. We also thank our colleagues in the MADAI (Models and Data Analysis Initiative) group (Chris Healey Cory Quammen, Olu-wafemi Alabi, and Hal Canary) for discussions and feedback. This research was supported partly by the US National Science Foundation's Cyber-Enabled Discovery and Innovation Program through grant 0941373 (madai.us), by Sandia National Laboratories contract 979573, and by US National Institutes of Health/National Institute of Biomedical Imaging and Bioengineering award 2P41EB002025 (cismm.org).

References

1. Harter JM, et al. Increasing the Perceptual Salience of Relationships in Parallel Coordinate Plots. Proc SPIE Visualization and Data Analysis 2012, Proc SPIE. 8294:82940T–82940T-12.
2. Cabral B, Leedom LC. Imaging Vector Fields Using Line Integral Convolution. Proc Siggraph. 1993:263–270.
3. van Wijk JJ. Spot Noise Texture Synthesis for Data Visualization. Proc Siggraph. 1991:309–318.
4. Urness T, et al. Effectively Visualizing Multi-valued Flow Data Using Color and Texture. Proc 14th IEEE Visualization Conf (VIS 03). 2003:115–122.
5. Urness T, et al. Strategies for the Visualization of Multiple 2D Vector Fields. IEEE Computer Graphics and Applications. 2006; 26(4):74–82. [PubMed: 16863100]
6. Interrante V, Grosch C. Strategies for Effectively Visualizing 3D Flow with Volume LIC. Proc 8th IEEE Conf Visualization (VIS 97). 1997:421ff.
7. Leech J, Taylor RM II. Interactive Modeling Using Particle Systems. Proc 2nd Discrete Element Methods Conf at MIT. 1992:105–116.
8. Ware, C. Information Visualization: Perception for Design. 3. Vol. ch 7. Morgan Kaufmann; 2012. Space Perception.

Biographies

Russell M. Taylor II is a research professor of computer science, physics, astronomy, and applied physical sciences at the University of North Carolina at Chapel Hill. Contact him at taylorr@cs.unc.edu.

Jonathan Harter is a staff scientist at BBN Technologies. He previously was a research assistant in the Models and Data Analysis Initiative at the University of North Carolina at Chapel Hill. Contact him at jmharter88@gmail.com.

Figure 1.

A hydrodynamic simulation of high-energy particle collisions, using (a) standard parallel coordinates and (b) luminance-modulated parallel coordinates. Adding random per-element luminance modulation improves the perception of detail in overplotted regions. In particular, the line bundles shaped like an S, a U, and an inverted U in the middle of the third parameter (U_{wind}) are more salient. Tracing an individual line across multiple axes is also easier.

Figure 2.

Animations of cones moving through a weather simulation (a) without and (b) with random luminance modulation. Without modulation, keeping track of a moving cone's neighbors is challenging. Modulation makes it easier to track individual cones and compare neighbors' trajectories. This is because there are fewer cones of the same color nearby that can become confused with each other during motion. Each cone thus becomes more perceptually distinct. (See the videos at <http://youtu.be/TTaSFMvBgvg> and <http://youtu.be/Rx1oPMTpPA4>.)

Figure 3. Streamline-based visualizations (a) without and (b) with random luminance modulation. Without modulation, the streamlines show overall flow well, but following a single tube's trajectory is challenging. Modulation makes it easier to trace individual tubes across long distances and compare neighbors' trajectories. In particular, the different luminance levels provide separation that helps viewers determine which of the tubes that enter the frame on the right edge exit the frame at the bottom, versus in the two bundles on the left. The relative depth of crossing lines is also enhanced.

Figure 4.

Isoluminant stress tensor hyperstreamlines from quantum-chromodynamics data, colored by gluon density, (a) without and (b) with random luminance modulation. With modulation, the density (hue) remains perceptible while the lines become distinct. You can see that the spiral near the top center consists of multiple distinct lines. Also, the lines making up the spiral near the right of center are highlighted against their neighbors as they wind around and out the bottom of the image.

Figure 5. Data-driven luminance variation in parallel coordinates.¹ This reveals the correlations between the first axis and each remaining axis. However, it doesn't consistently separate local clusters of lines (see the upper halves of T_sonic and soil_H2O).