

Coloring 3D Line Fields Using Boy's Real Projective Plane Immersion

Çağatay Demiralp, John F. Hughes, and David H. Laidlaw, *Senior Member, IEEE*

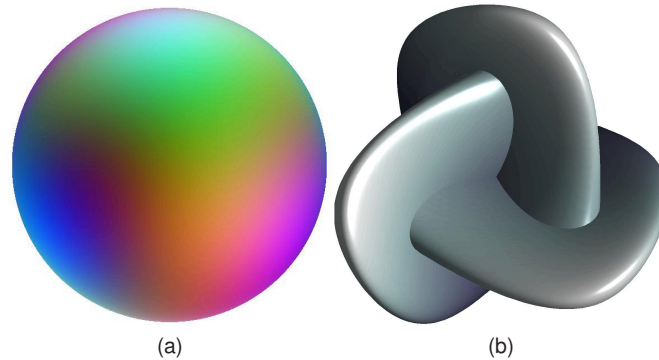


Fig. 1: a) Sphere colored by immersing RP^2 in RGB color space b) Boy's surface

Abstract—We introduce a new method for coloring 3D line fields and show results from its application in visualizing orientation in DTI brain data sets. The method uses Boy's surface, an immersion of RP^2 in 3D. This coloring method is smooth and one-to-one except on a set of measure zero, the double curve of Boy's surface.

Index Terms—Line field, colormapping, orientation, real projective plane, tensor field, DTI.

1 INTRODUCTION

It is often useful to visualize a *line field*, a function that sends each point P of the plane or of space to a line through P (see Figure 2a-b); such fields arise, for example, in the study of tensor fields, where the principal eigendirection at each point determines a line. The direction is a line rather than a vector because diffusion is spatially symmetric; if \mathbf{v} is an eigenvector, so is $-\mathbf{v}$. Second-order tensor field data is common in physics, engineering, medical imaging, and structural imaging. Line fields also appear in computer vision, material science, and chemistry as edge, orientation, or rotation fields.

To visualize a line field, we often assign a color to each line, typically according to its direction. This can be thought of as a mapping from the set of all line-directions to color space. When we color a line field, there are some natural desiderata:

- **Smoothness:** the color should change smoothly as the line changes smoothly; i.e., the color-triple should be a smooth function of the coordinates of the direction-vector of the line;
- **Locality:** lines with similar directions should map to similar colors;
- **Uniqueness:** lines with very different directions should map to very different colors.

Previous methods for line field coloring suffer from either discontinuities or uniqueness problems (i.e., multiple lines corresponding to a single color). One reason is that the problem, as formulated, has no solution! Finding a mapping from RP^2 to the space of color triples, which is topologically a subset of \mathbf{R}^3 , while satisfying the above criteria is equivalent to embedding RP^2 in 3-space. But it is possible to

show that no such mapping exists [9]: RP^2 is a nonorientable surface, it admits no embedding in 3-space.

Contributions Our primary contribution is introducing Boy's RP^2 immersion as a color model for line fields. The resulting coloring method is smooth and one-to-one except on a set of measure zero. We also present an unpublished parametrization of Boy's surface that was found 15 years ago by one of the authors.

We apply the proposed model to color the orientation in slices of diffusion tensor magnetic resonance imaging (DTI) brain volumes (a second-order tensor field) and integral curves, which represent the fiber tracts, extracted from these volumes. DTI provides a non-invasive means to measure microarchitecture of fibrous tissues induced by the Brownian motion of water molecules [3]. Compare to conventional MRI, DTI enables the acquisition of much richer anatomical information about the white matter structures (myelinated axons) in the brain. During the last decade this has led to a number of important applications for clinical and neuroscience research into the health of the brain [11].

1.1 Related work

Color-mapping of line fields has been studied in the context of diffusion tensor magnetic resonance imaging (DTI) visualization widely. There have been several methods based on rotational and mirror-symmetry assumptions. An extensive study of different coloring methods for visualizing DTI can be found in [13]. One standard method for visualizing DTI slices is to directly map the absolute values of the xyz -coordinates of principal eigenvectors to RGB color-triples [14]. We refer to this method as the *absolute value method* throughout this paper. Clearly, this does not have uniqueness: there are 4 different orientations associated to the color (r, g, b) whenever $r, g,$ and b are all nonzero (Figure 2-c).

However, the absolute value method provides an easy-to-read color model, as it aligns the major colors with the major axes. As the data also tend to be aligned with these major axes, colors can often be compared across datasets. There have been other methods proposed in dif-

• Çağatay Demiralp, John F. Hughes, and David H. Laidlaw are with Brown University, E-mail: {cad,jfh,dhl}@cs.brown.edu.

Manuscript received 31 March 2009; accepted 27 July 2009; posted online 11 October 2009; mailed on 5 October 2009.

For information on obtaining reprints of this article, please send email to: tvcg@computer.org.

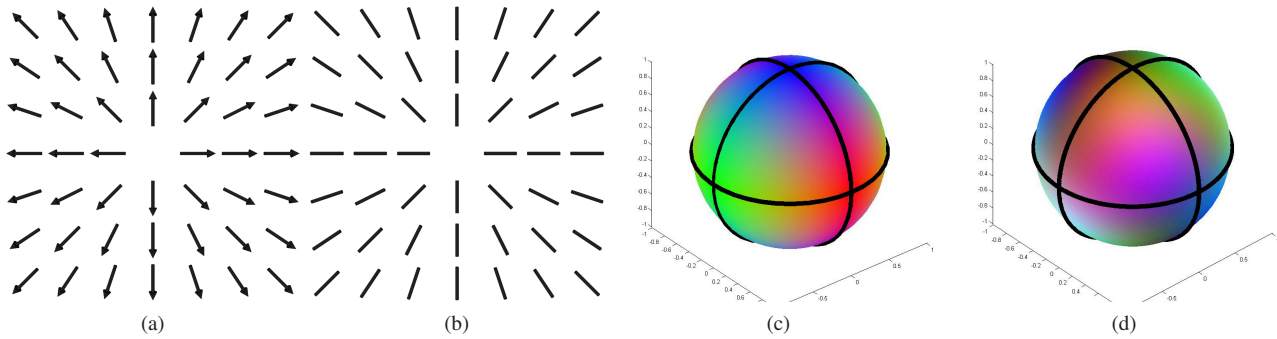


Fig. 2: a) A vector field. b) A line field. Sphere colored using c) absolute value method and d) Boy's immersion. Three great circles aligned with the major axes are also shown. Note that in the case of the absolute value method colors are repeated 8 times, once within each spherical triangle bounded by the great circles.

ferent contexts such as visualization of material microstructures [16] and rotation fields; since every nonzero 3D rotation has a unique axis of rotation a rotation field induces a line field [1]. All these methods, however, suffer either from uniqueness or continuity problems. The coloring method proposed here is continuous and one-to-one almost always. Furthermore, in the few cases where uniqueness is not satisfied, context usually provides disambiguation.

There have been different parametrizations proposed for Boy's surface. The first parametrization was found by Morin to be used as a half-way model for sphere eversion [12, 7]. Based on Morin's work, his student Apery found another parametrization [2, 7]. Bryant and Kusner also proposed a smooth immersion which minimizes the Willmore energy of the surface (the least bent and most spherical) using complex arithmetic [5, 8]. Our approach is to represent a Boy-like immersion using low-degree polynomials restricted to the unit sphere, determined by treating Boy's immersion as an approximation to a polyhedral shape.

2 THE REAL PROJECTIVE PLANE

The real projective plane (RP^2) is the set of lines through the origin in R^3 ; this set corresponds naturally to a surface. There are several ways to describe this surface; we show the most common ones in Figure 3.




Fig. 3: Models of RP^2 . Top: The upper hemisphere, with antipodal points on the equator identified. Middle: a unit disk, with antipodal boundary points identified. Bottom: A rectangle, with opposite edges identified, with a reversal in orientation.

The fundamental idea in all our descriptions is that each element of RP^2 — each line through the origin — intersects the upper hemisphere somewhere. For most lines, the intersection point is unique. But for horizontal lines, there are two intersections, at antipodal points on the equator. This establishes a correspondence between RP^2 and a set whose elements are (a) points in the interior of the upper hemisphere and (b) pairs of antipodal points on the equator. If we imagine “gluing together” antipodal points (just as we “glue together” the two copies of the international dateline in wrapping a mercator projection onto a sphere), we get a representation of RP^2 ; unfortunately, it's hard to imagine doing this in three dimensions.

The upper hemisphere is topologically the same as a disk, so we can instead consider RP^2 as a disk with antipodal boundary points identified. And a disk is topologically equivalent to a rectangle, so we can consider RP^2 as a rectangle with opposite sides identified. All three of these approaches are shown schematically in Figure 3.

2.1 Realization in 3D and Boy's Surface

Because the real projective plane is a non-orientable surface (i.e., it contains a Möbius strip) and is closed (has no boundary), it admits no embedding in R^3 . This implies that there can be no one-to-one and smooth coloring of 3D line fields.

Nonetheless, it is still possible to immerse RP^2 in 3-space in a relatively simple way. A smooth mapping $f : M \rightarrow R^n$ is an immersion of the surface M if $Df(p)$ the derivative at the point p , has rank two at each point $p \in M$, and if for every point $p \in M$, there is a neighborhood U_p of p such that f , restricted to U_p , is one-to-one; intuitively speaking, an immersion is a mapping that looks like an embedding locally, but which may have self-intersections. The first immersion of RP^2 in 3-space was discovered by the German mathematician Werner Boy using hand-drawn figures in his 1901 thesis [4]; this immersion is known as *Boy's surface*. It is this three-fold symmetric immersion of RP^2 that we use for coloring line fields.

2.2 Parametrization

Boy's immersion was described visually; we generated a smooth mapping from RP^2 to R^3 whose image has the general shape indicated by Boy. We first generated a polyhedral shape (the *boxy approximation*) that approximated Boy's immersion, in much the way that a cube could be said to approximate a sphere (see Figure 4). We then triangulated the unit sphere with an antipodally-symmetric triangulation (i.e., if P is a vertex, so is $-P$; if there is an edge from P to Q , there is an edge from $-P$ to $-Q$, etc.) so that each antipodal pair of vertices in the triangulation corresponded to a single point in the boxy approximation, and each antipodal pair of triangles on the sphere corresponded to a single triangle in the boxy approximation




Fig. 4: Boxy-approximation of Boy’s surface, shown here as a union of three congruent pieces in a) partially-exploded, b) intermediate, and c) assembled views.

of Boy’s surface. This defines a mapping f from the sphere to the boxy approximation: each point $\pm(x, y, z)$ in a spherical triangle was represented by barycentric coordinates with respect to its spherical triangle; it was mapped to a point in the corresponding triangle in the boxy approximation with the same barycentric coordinates. We write $f(x, y, z) = (f_1(x, y, z), f_2(x, y, z), f_3(x, y, z))$.

The continuous function $f : S^2 \rightarrow \mathbf{R}^3$ could be described as a “piecewise linear version of Boy’s immersion.” Each of the three component functions f_i of f is a continuous real-valued function on the sphere. Every such function can be written as an infinite weighted sum of spherical harmonics, using a method completely analogous to Fourier analysis, which expresses a periodic function on the real line as a weighted sum of sines and cosines [6].

The first component function, f_1 , therefore can be written in the form

$$f_1(x, y, z) = \sum_{i=0}^{\infty} c_i h_i(x, y, z)$$

where the h_i are the spherical harmonics, and the c_i are real-valued weights. Because of the antipodal symmetry ($f_1(x, y, z) = f_1(-x, -y, -z)$ for all x, y, z), only the spherical harmonics of even total degree (when expressed as polynomials in x, y , and z) can have nonzero coefficients in this expansion.

As defined, f_1 is continuous but not smooth. To smooth it, we can take a *partial* sum:

$$f_1(x, y, z) \approx \sum_{i=0}^n c_i h_i(x, y, z)$$

of just n terms of its “Fourier series”. Since each spherical harmonic is smooth, the resulting function is a smooth approximation of f_1 . We can do the same with f_2 and f_3 , of course.

This is similar to Fourier-analyzing a square-wave s (which is discontinuous) as

$$s(x) = \sin(x) + \frac{1}{3}\sin(3x) + \frac{1}{5}\sin(5x) + \dots$$

and then noting that the first term of the series gives a smooth approximation of the discontinuous square wave. If we sum the first several terms, to improve the approximation, we get “ringing artifacts” in which the approximation “overshoots” the square wave near discontinuities. If we instead sum the first n terms, but multiply the i th

coefficient by $\frac{n-i}{n}$, this artifact is substantially reduced; this is called “smoothing with the Fejer kernel.”

We did the same thing with Boy’s surface: we ordered the spherical harmonics by degree (each spherical harmonic can be written as a polynomial in the coordinates x, y , and z on S^2), sum up the low-degree terms in the Fourier series, attenuating by degree; the result is a smooth approximation of the boxy approximation; the smooth approximation is expressed as three fourth-degree polynomials in the x, y, z coordinates on the unit sphere. As a final step, we rounded the coefficients to two decimal places, and then adjusted some coefficients by hand to achieve an aesthetically pleasing shape, because our original goal (25 years ago) in creating this polynomial immersion was aesthetically motivated. The entire process of finding this polynomial immersion was ad hoc: the boxy approximation was easy to create; the triangulation of the boxy faces was done to maximize minimum angles; the spherical triangulation was hand-adjusted to have “nicely shaped” triangles whose areas were similar to those of the corresponding triangles in the boxy approximation.

3 COLORING

Given the parametrization above (or any other parametrization of Boy’s surface), coloring is straightforward. To color an element of a line field (which is determined, up to sign, by a unit vector $\mathbf{v} = (x, y, z)$), we consider the element $[x, y, z] \in RP^2$; to this, we apply Boy’s immersion f , producing an element $f([x, y, z]) \in \mathbf{R}^3$, which we in turn treat as a color-triple. We implemented our coloring based on Boy’s immersion in Matlab and make it available online at [15].

3.1 Color Space of Immersion

We immerse Boy’s surface both in the RGB and L*a*b* color spaces. L*a*b* is a perceptually uniform color space, where the perceptual difference between any two colors is approximately equal to the Euclidean distance between them in the color space. For example, if there is an isometric map from data points to a Euclidean subspace, the L*a*b* would provide a coloring reflecting “distances” between the data points as perceptual differences. Since Boy’s immersion is smooth and one-to-one locally, the L*a*b* is a natural choice. However, the gamut of the L*a*b* color space has an irregular shape, and saturated colors lie close to the gamut boundary. This consequently limits the utilization of the saturated colors. On the other hand, the RGB has a rectilinear shape allowing simpler use of the color volume. Note also that, in both color spaces, the level sets of luminance are




Fig. 5: Tracts colored by encoding the lines determined by the end-points of the tracts (also known as the *end-point vector*, direction of which does not matter); a common way coloring fiber tracts in DTI. a) Synthetic tracts colored using (top) the absolute value method and (bottom) Boy’s surface immersion in RGB; Tracts designed to demonstrate the symmetries of the absolute value method. The internal capsule, extracted from a normal person’s DTI brain data set, colored by encoding the end-point vector with b) the absolute value method, and c) Boy’s surface. While Boy’s RP^2 immersion assigns unique colors to unique lines, the absolute value method assigns the same color to the mirror-symmetric lines.

two-dimensional planes; Therefore, it is not possible to create an iso-luminant color mapping using a non-planar immersion like Boy’s.

4 RESULTS

We applied our model to sectional visualizations of DTI brain volumes and to fiber tracts extracted from these volumes by tracking principal eigenvectors of the underlying tensor-field. For comparison, we also obtained coloring results for the standard absolute value method. The DTI datasets were gathered by scanning two healthy subjects. Results, shown in Figures 5-9, demonstrate that coloring based on Boy’s surface assigns unique colors to unique tissue orientations in a smooth, continuous manner, while the absolute value method does not.

5 DISCUSSION

Note that Boy’s surface has a curve of self-intersections that winds around and ultimately intersects itself at a triple point. This means that certain pairs of distinct directions will be coded by the same color. But, unlike the absolute value coloring method, directions nearby these will be coded differently, so surrounding context should often help distinguish the directions. Also, it should be possible to obtain a solution of the double curve numerically and ϵ -perturb the points on it.

The power of our coloring scheme (or any other smooth, nearly-unique continuous mapping) shows itself not only in visualization but also in data inspection for verification (e.g., see Figure 9). RP^2 is a natural model for studying line fields and has been used in physics, particularly for describing singularities in nematic crystals topologically [10]. We expect it to find similar applications in the analysis and visualization of tensor fields in the future.

5.1 Comparing Boy’s Surface with the Absolute Value Method

While uniqueness and smoothness are important criteria, they are not the only ones. For example, the absolute value method is an easy-to-read color model because it aligns the additive color primaries with the coordinate directions of the DTI datasets. Our parameterization of Boy’s surface happens to send six points fairly near the equator (i.e., three line-pairs that pierce the sphere near the equator) to the tips of the three lobes of the immersion, while the north/south-pole pair is sent to the “center of the cap”; it is possible, by applying a rotation to the sphere, and adjusting the triangulation slightly, to make a very similar parameterization that sends the directions of the three major axes to the tips of the three lobes, which then can be placed on the three axes in color space, so that the Boy’s-surface-based coloring shares the easy-to-read property with the absolute value method, at least for directions near the primary directions; we have not implemented this, however. We instead show coloring results for an immersion of the Roman surface transformed to maximize the alignment with major colors in the following section.

The absolute value method has a four-fold symmetry and, hence, can be useful for displaying structures with four-fold symmetry. One of the reasons that the absolute value method has been popular in visualizing DTI brain data sets is the fact that the brain has a two-fold (left-right) symmetry across the midsagittal plane at a certain scale. However, such symmetry does not exist in other directions, or even in the left-right direction within the hemispheres of the brain, not to mention other organs of interest such as the heart. Furthermore, it is always possible to turn a one-to-one color mapping into a symmetric one, given a reference of symmetry. Therefore it is a color model with




Fig. 6: Mid-coronal slice sagittal slice from a DTI brain data set of a normal subject; colored by mapping principal tensor axis at each point to an RGB color using a) the absolute value method and b) Boy's RP^2 immersion. Major anatomical structures selected, including the corpus callosum and internal capsule, show that while Boy's immersion assigns different colors to different tissue orientations, the absolute value method does not.




Fig. 7: Left-sagittal slice from a DTI brain data set of a normal subject; colored by mapping principal tensor axis at each point to an RGB color using a) the absolute value method and b) Boy's RP^2 immersion. Major anatomical structures selected include the cingulum bundle and corpus callosum.




Fig. 8: The internal capsule and the middle cerebellar peduncle (the tilted U-shape structure). Tracts colored by encoding the end-point vector (direction of which does not matter) with a) the absolute value method, b) Boy's surface immersed in RGB, and c) Boy's surface immersed in $L^*a^*b^*$, an approximation for a perceptually uniform color space.




Fig. 9: a) Fiber tracts obtained from the whole DTI brain dataset of a healthy subject. Tracts colored encoding the end-point vector via immersing Boy's surface in RGB. b) Unique coloring helps in this case finding c) broken curves (along the yellow line).




Fig. 10: Sphere colored by immersing different realizations of RP^2 in RGB color space. a) Bryant-Kusner parametrization of Boy's surface, b) Cross-cap, c) Roman surface.

the uniqueness property would typically be preferable to one without.

The absolute value method “squeezes” the color space into a spherical triangle with area one eighth of the sphere. The eight identically colored copies of the triangle are shown in Figure 2-c. Because the vertices of this triangle are fully saturated, all of the colors tend to be quite saturated, which can be desirable in certain application settings. This also hints of a trade-off between the color saturation and “bandwidth.” Both of the methods can be improved, in the latter sense, by increasing the surface area within the color space by superimposing smooth wrinkles.

5.2 Other Realizations of RP^2 in R^3

Boy's surface has other well-known parametrizations as we discussed above. A particularly elegant one is the Bryant-Kusner parametrization. This parametrization is the optimum mapping from RP^2 to R^3 in the sense that it minimizes the bending (or Willmore) energy, which is equivalent to locally maximizing sphericity. This feature of the Bryant-Kusner parametrization can help conveying dissimilarities more faithfully, particularly when a perceptually uniform color space is used. There are also other realizations of RP^2 that were known before Boy's surface, including the cross-cap and Roman surface, both of which are Steiner surfaces with singularities. We show the surface models of these realizations and coloring results obtained by immersing them in the RGB color space in Figure 10.

Which coloring is perceptually optimum entails further investigation, possibly including a user study. While both parametrizations of Boy's surface have nice properties such as non-singularity, smoothness, minimal distortion, and better coverage area, the Roman surface or cross-cap may be easier to align with the major axes. For example, by applying simple transformations (rigid and non-rigid) to the Roman surface, we obtained colorings approximately aligned with the major color axes (see Figure 11). As pointed in the previous section, it is possible to construct a similar alignment for Boy's surface as well, combining its desirable properties with ease of use.

6 SUMMARY AND CONCLUSIONS

Any coloring of a 3D line field is a mapping from RP^2 to a Euclidean subspace. We have introduced Boy's immersion of RP^2 as a method for such a coloring. Given that a smooth, one-to-one coloring for line fields is impossible in general, the coloring based on Boy's surface is a

good approximation, because it is one-to-one and smooth except along the self intersection curve of Boy's surface. While we have applied our model to visualization of DTI-derived data, it is likely to be useful in any domain where data instances are represented as lines, including edge, orientation, and rotation fields.

ACKNOWLEDGMENTS

Çağatay Demiralp thanks Thomas Banchoff for insightful discussions on geometry and topology. Also, John Hughes thanks Brian Cole for suggesting smoothing via the Fejer kernel. This work has been supported partly by NIH grant number 1R01EB00415501A1 and the Reisman Fellowship of the Brown Institute for Brain Sciences.

REFERENCES

- [1] B. Alpern, L. Carter, M. Grayson, and C. Pelkie. Orientation maps: techniques for visualizing rotations (a consumer's guide). In *VIS '93: Proceedings of the 4th conference on Visualization '93*, pages 183–188, 1993.
- [2] F. Apery. *Models of the Real Projective Plane*. Friedrich Vieweg & Sohn Verlag, 1987.
- [3] P. J. Basser, J. Mattiello, and D. LeBihan. Estimation of the effective self-diffusion tensor from the nmr spin echo. *J Magn Reson B*, 103(3):247–254, March 1994.
- [4] W. Boy. *Über die Curvatura Integra und die Topologie der Geschlossener Flächen*. PhD thesis, Universität Göttingen, Göttingen, 1901.
- [5] R. L. Bryant. Surfaces in conformal geometry. In *The mathematical heritage of Hermann Weyl (Durham, NC, 1987)*, volume 48 of *Proc. Sympos. Pure Math.*, pages 227–240. Amer. Math. Soc., Providence, RI, 1988.
- [6] W. E. Byerly. *An Elementary Treatise on Fourier's Series, and Spherical, Cylindrical, and Ellipsoidal Harmonics, with Applications to Problems in Mathematical Physics*. Dover, New York, 1959.
- [7] G. K. Francis. *A Topological Picture Book*. Springer, 2006.
- [8] R. Kusner. Conformal geometry and complete minimal surfaces. *Bull. Amer. Math. Soc. (N.S.)*, 17(2):291–295, 1987.
- [9] W. S. Massey. *Algebraic Topology: An Introduction*. Number 56 in Graduate Texts in Mathematics. Springer-Verlag, New York, NY, 1967.
- [10] V. Mineev. Topological objects in nematic liquid crystals. In *Intuitive Combinatorial Topology*. Springer, July 2001.
- [11] S. Mori and J. Zhang. Principles of diffusion tensor imaging and its applications to basic neuroscience research. *Neuron*, 51:527–539, 2006.
- [12] B. Morin. Équations du retournement de la sphère. *C. R. Acad. Sci. Paris Sér. A-B*, 287(13):A879–A882, 1978.
- [13] S. Pajevic and C. Pierpaoli. Color schemes to represent the orientation of anisotropic tissues from diffusion tensor data: Application to white matter fiber tract mapping in the human brain. *MRM*, 42:526–540, 1999.
- [14] C. Pierpaoli. Oh no! one more method for color mapping of fiber tract direction using diffusion mr imaging data. In *Proceedings of 5th ISMRM*, page 1741, 1997.
- [15] <http://www.cs.brown.edu/~cad/rp2coloring/>.
- [16] B. Yamrom, J. A. Sutliff, and A. P. Woodfield. Visualizing polycrystalline orientation microstructures with spherical color maps. In *VIS '94: Proceedings of the conference on Visualization '94*, pages 46–51, Los Alamitos, CA, USA, 1994. IEEE Computer Society Press.




Fig. 11: Sphere colored by a) using the absolute value method, b) immersing the Roman surface rotated to align with the major color axes in RGB, and c) immersing the Roman surface transformed (non-rigid) to align with the major color axes in RGB.