4-2016

Polyhedral Painting with Group Averaging

Frank A. Farris

Santa Clara University, ffarris@scu.edu

Ryan Tsao

Follow this and additional works at: http://scholarcommons.scu.edu/math_compsci

Part of the [Mathematics Commons](#)

Recommended Citation

Farris, Frank A., and Ryan Tsao. "Polyhedral Painting with Group Averaging." *Primus* 26.4 (April 2016): 312-24.

This is an Accepted Manuscript of an article published by Taylor & Francis in *PRIMUS: Problems, Resources, and Issues in Mathematics Undergraduate Studies* in v. 26 no. 4 in April 2016, available online: <http://www.tandfonline.com/10.1080/10511970.2015.1110218>.

This Article is brought to you for free and open access by the College of Arts & Sciences at Scholar Commons. It has been accepted for inclusion in Mathematics and Computer Science by an authorized administrator of Scholar Commons. For more information, please contact rscroggin@scu.edu.

PRIMUS

Problems, Resources, and Issues in Mathematics Undergraduate Studies

ISSN: 1051-1970 (Print) 1935-4053 (Online) Journal homepage: <http://www.tandfonline.com/loi/upri20>

Polyhedral Painting with Group Averaging

Frank A. Farris & Ryan Tsao

To cite this article: Frank A. Farris & Ryan Tsao (2015): Polyhedral Painting with Group Averaging, PRIMUS, DOI: [10.1080/10511970.2015.1110218](https://doi.org/10.1080/10511970.2015.1110218)

To link to this article: <http://dx.doi.org/10.1080/10511970.2015.1110218>

Accepted author version posted online: 29 Dec 2015.

Submit your article to this journal [↗](#)

View related articles [↗](#)

View Crossmark data [↗](#)

Polyhedral Painting with Group Averaging

Frank A. Farris

Department of Mathematics and Computer Science Santa Clara University Santa Clara, CA
95053 ffarris@scu.edu 408-554-4430

Ryan Tsao

Department of Mathematics and Computer Science Santa Clara University Santa Clara, CA
95053 ryan.j.tsao@gmail.com

Polyhedral Painting with Group Averaging

Abstract: The technique of *group averaging* produces colorings of the sphere that have the symmetries of various polyhedra. The concepts are accessible at the undergraduate level, without being well known in typical courses on algebra or geometry. The material makes an excellent discovery project, especially for students with some background in computer science; indeed, this is where the authors first worked through the material, as teacher and student, producing a previously unseen type of artistic image. The process uses a photograph as a palette, whose colors and textures appear in kaleidoscopic form on the surface of a sphere. We depict tetrahedral, octahedral, and icosahedral symmetries, with and without mirrors, along with the source photograph for comparison. We also describe a method to make images with color-reversing symmetry.

Keywords: symmetry, group, polyhedra, color-reversing, mathematical art

1 FUNCTIONS AND SYMMETRY

Most treatments of plane symmetry [1, 4] describe patterns as being built from many copies of a *motif*, a colored swatch of the plane that is repeated to fill the plane without gaps or overlaps. Recent work by Farris, culminating in [2], finds mathematical and artistic interest in describing patterns as real- or complex-valued functions, whether on the plane or other spaces. Indeed a chapter in [2] treats polyhedral symmetry, but the techniques there—Fourier series and complex analysis—are different from what we offer here.

Figure 1. A polyhedral painting with tetrahedral (and mirror) symmetry. The colors come from a photograph of an octopus, shown in Figure 4.

As a project for Tsao and another student, Matthew Medal, in a course called Survey of Geometry, Farris proposed an approach not involving complex numbers, which led to new understandings, as well as images such as Figure 1.

We start with a real-valued function on $\mathbb{R}$ as a simple example to construct symmetry in the context of functions. The simplest symmetry such a function f can have is to be an *even function*, which means that

$$f(-x) = f(x).$$

We usually visualize this symmetry in terms of the graph of f : The function f is even if and only if its graph has reflection symmetry across the y -axis. How might one construct an even function? Many calculus students have seen the formula

$$f_e(x) = \frac{f(x) + f(-x)}{2}, \quad (1)$$

which creates an even function f_e from an arbitrary base function f . The term *even* becomes vivid when we realize that, if f is defined by a power series, then f_e consists exactly of the terms in that series where the power of x is even. Of course, if we chose a function f whose power series contains only odd powers, the even function in (1) would turn out to be zero.

This simple construction technique gives us a basic example of *group averaging*. Let us name the reflection map

$$R(x) = -x \text{ and call the identity map } I(x) = x.$$

Then the set $G = \{I, R\}$ fits the definition of a *group* under composition: It is closed under composition, it contains an identity elements, the operation is associative, and, since $R^2 = I$, it contains the inverse of each of its elements. An even function is nothing more than one that is invariant under the action of every element of the group: Every function is invariant under I and invariance under R is the defining property of an even function.

We call the process in (1) *group averaging* because we have summed the compositions of f with each element of the group and divided by the number of elements in the group. The resulting average is evidently invariant under the action of every element of the group.

To generalize this idea, any time we have a finite group G with $|G|$ elements that acts as transformations of a set X , then we define the *average of f over G* to be

$$\hat{f}(x) = \frac{1}{|G|} \sum_{g \in G} f(g(x)) \text{ for } x \in X. \quad (2)$$

A simple check shows that

$$\hat{f}(g(x)) = \hat{f}(x) \text{ for every } g \in G, x \in X.$$

Group averaging does indeed produce invariant functions. As we observed for even functions, the averaged function may turn out to be zero if the base function is chosen poorly.

2 TETRAHEDRAL SYMMETRY

Many students in an algebra class know that the symmetries of a tetrahedron, not allowing reflection symmetry, constitute the group A_4 , the alternating group on four elements. This is because we can find a rotation in space that permutes the four vertices of the tetrahedron in every possible way, as long as the permutation is even. However, students sometimes skip the concreteness that might come from realizing these symmetries as transformations.

In fact, the tetrahedral group may be generated by three elements. Let us define H_x to be the 180° rotation about the x -axis, H_z similarly about the z axis, and call the cyclic permutation of the axes, which is a 120° rotation about the vector $(1, 1, 1)$, P . In matrices, we write

$$H_x = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & -1 \end{pmatrix}, H_z = \begin{pmatrix} -1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \text{ and } P = \begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}.$$

Figure 2 is meant to help you see that these transformations preserve the tetrahedron with vertices $(1, 1, 1)$, $(-1, -1, 1)$, $(1, -1, -1)$, and $(-1, 1, -1)$. Two views show the tetrahedron sitting inside a cube, with the second view showing why P is a symmetry.

Figure 2. Two views of a tetradron sitting inside a transparent cube.

It is fun to see how these fit together to form the tetrahedral group. First we notice the subgroup of rotations about the vector $(1, 1, 1)$. Let us call it

$$S = \{I, P, P^2\},$$

where I is the identity transformation. The full tetrahedral group is the union of four (right) cosets:

$$T = S \cup SH_x \cup SH_z \cup SH_xH_z.$$

(A formal proof requires us to verify several things, for instance, that the composition H_xH_z is a half turn about the y -axis, which we call H_y , that $PH_x = H_yP$, and so on.)

Note that T contains only the rigid motions of the tetrahedron, without allowing reflection symmetries. For this reason, some call it the *chiral* tetrahedral group, where the word *chiral* describes objects lacking mirror symmetry. Later, we will study possibilities for including mirror symmetry.

To apply group averaging and create our first polyhedral painting, we start with a function on the sphere. We could use an ordinary real-valued function, but we found greater interest in *color-valued* functions on the sphere. By this we mean

$$f : S^2 \rightarrow \mathbf{C} = \{(R, G, B) \mid R, G, \text{ and } B \text{ are integers from } 0 \text{ to } 255\}.$$

After all, colors on a computer monitor are specified by three integers in the given range, specifying the amounts of Red, Green, and Blue. We freely add colors just as if they were vectors, which you might think would lead to breaking the range restriction. One solution to this difficulty would be to add colors modulo 256, creating discontinuous color breaks in images. However, since we are *averaging* values, they will stay nicely within the desired range.

Therefore, if we have any color-valued function on the sphere, f , we can form

$$\hat{f}(\vec{x}) = \frac{1}{12} \sum_{g \in T} f(g\vec{x})$$

for $\vec{x} \in S^2$, and find that $\hat{f}$ is invariant under the tetrahedral group.

Now all we need is an initial color-valued function on the sphere. Really, any way to assign colors to points on the sphere will work, but we used photographs in conjunction with ordinary spherical coordinates. Starting with a photo with an aspect ratio of 2 to 1, we map the long dimension to the angular variable around the sphere—the longitude—and the short dimension to the angle down from the pole, a shifted version of latitude. If the photograph has $2L$ by L pixels,

the point in the photo with pixel coordinates (A, B) (A measured left to right and B measured top to bottom) ends up at the point

$$\left(\cos\left(\frac{\pi A}{L} + \frac{\pi}{4}\right) \sin(\pi B / L), \sin\left(\frac{\pi A}{L} + \frac{\pi}{4}\right) \sin(\pi B / L), \cos(\pi B / L) \right)$$

on the unit sphere. (We have turned the sphere by 45° relative to the usual conventions. Doing so gives our mirror formulas a nicer appearance.) This creates discontinuities in the image on the sphere, where the left-hand edge of the photo meets the right-hand edge along a meridian. It also distorts the photo at the poles, cramming the top and bottom edges of the picture onto the top and bottom poles. Knowing this affects our choice of photographs. More on this point later.

Tsao and Medal used WebGL to realize this routine for putting a photograph onto a sphere and then averaging it over the tetrahedral group. The software allows the user to choose various photographs and rotate the painted sphere in order to see its features better. After all, it is difficult to confirm tetrahedral symmetry when one can only see one side of the sphere.

The program also allows the user to shift the photograph left or right, up or down, wrapping the edges around, as if it were a colored torus. A slider to control “X-shift” allows us to use the colors at the point $((A + X) \bmod 2L, B)$ instead of those at (A, B) . A “Y-shift” slider works similarly. The discontinuities of color produced by these shifts can create interesting features in the result.

An example appears in Figure 3, where we can see the original photograph, the way it was mapped onto the sphere, and the result after averaging, with this *caveat*: If the source image uses

widely contrasting colors, averaging desaturates them, moving them towards the center of the color cube—a dull gray (128, 128, 128) as (R, G, B) . To maintain contrast for aesthetic purposes, we post-process the image by boosting saturation.

Figure 3. A photograph of a sunset, the same image mapped onto the sphere, and the average of this photograph over the tetrahedral group T . (Image has been post-processed to increase color saturation.)

To help the reader see the symmetry, we mention that there are two different centers of 3-fold rotational symmetry visible on the left- and right-hand sides of the sphere as displayed. One corresponds to a vertex, the other to a face of the tetrahedron. The centers of 2-fold rotational symmetry at the top and bottom of the image correspond to centers of the tetrahedron's edges.

Discontinuity of the original painting on the sphere contributes to the striping effect. Here, we think it works nicely; in other images, one may prefer that the colors at top and bottom, left and

right of the image match more closely. The question of which photographs, or photographic collages, contribute best to a beautiful result after group averaging is a completely open question for artists to explore.

3 MIRROR SYMMETRY AND COSET AVERAGING

You might have noticed that Figure 1 has more symmetry than we construct in the previous section: mirror symmetry. That image inherited its mirror symmetry from its source image, shown in Figure 4. We used Photoshop to paste a reflected image of the octopus (taken in the National Aquarium in Baltimore), creating mirror symmetry about a vertical axis.

Figure 4. A source image of an octopus, rigged to exhibit mirror symmetry, and the sphere colored by that image. It is somewhat surprising that this symmetry survives group averaging, as we see in Figure 1.

When we transfer this specially arranged image onto the sphere, the resulting painting is invariant under the reflection across the plane $x = y$, which we call the *diagonal mirror*, as the plane of reflection contains the long diagonal of the cube that is the axis of that 120° rotation.

We label this reflection as

$$M_d = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

(The simplicity of this formula is why we chose the offset of $\pi/4$ in our photo-to-sphere mapping.)

When we use this image to color the sphere and then average over the tetrahedral group, the averaged function retains the mirror symmetry of the source! The reason for this is a variation on the group-averaging technique, called *coset averaging*. Let's develop that principle in the abstract, before we apply it to the mirror symmetry of the tetrahedron.

If H is any subgroup of a group G , then we call the set $\{g_1, g_2, \dots, g_k\}$ a complete set of *coset representatives for H in G* if the group G can be written as a disjoint union

$$G = H \cup Hg_1 \cup Hg_2 \cup \dots \cup Hg_k,$$

where Hg_j denotes the set $\{hg_j \mid h \in H\}$. For example, in our discussion of $\mathbb{T}$ above, the set $\{I, H_x, H_z, H_x H_y\}$ is a complete set of coset representatives for the subgroup S in $\mathbb{T}$.

It is a good exercise to verify that if we start with a function invariant under the subgroup H ,

$$f(\vec{x}) \text{ such that } f(h\vec{x}) = f(\vec{x}),$$

then the new function

$$\hat{f}(\vec{x}) = \sum_{j=1}^k f(g_j \vec{x}),$$

where $\{g_1, g_2, \dots, g_k\}$ is a complete set of coset representatives for H in G , is invariant under *all* elements of G .

The reason our trick with the octopus worked is that the specially-contrived image, when pasted onto the sphere by our algorithm, creates a color function f that is invariant under the 2-element group

$$\mathbb{M}_d = \{I, M_d\}.$$

The elements of the tetrahedral group T are in fact a set of coset representatives for the supgroup M_d in T_d , the full set of isometries of the tetrahedron, not just the rigid motions. We define this group by

$$T_d = \langle T, M_d \rangle,$$

meaning the smallest group containing T and M_d . When we use group averaging for the special image, we are actually doing coset averaging for the larger group. A neat trick.

There is another group obtained by T and a reflection. Let's name

$$M_{xy} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -1 \end{pmatrix},$$

the reflection about the xy -plane. Refer back to Figure 2 to see that this reflection preserves the cube in which the tetrahedron sits, but flips the tetrahedron drawn in that figure with a different tetrahedron in the same cube. One name for this group (in a system of naming called *Schoenflies notation*) uses the letter h as a mnemonic for this *horizontal mirror*:

$$T_h = \langle T, M_{xy} \rangle.$$

Our trick works just the same to create images invariant under T_h ; we just rig the source image to have symmetry across a horizontal axis. Figure 5 shows a painting with T_h symmetry, along with the source image we used to produce it. Notice the rotor effect at the centers of 3-

fold rotational symmetry. The lack of mirror planes through these points is a distinctive feature of this type of symmetry.

4 OTHER POLYHEDRAL GROUPS

Beyond the three groups we have used to paint the sphere, T , T_d , and T_h , there are only four more possibilities, and these involve octahedral and icosahedral symmetry, with and without mirrors. It turns out that there are only seven groups of symmetries of the sphere with the property that no axis is fixed by all elements of the group. These are the *point groups* of the 3D crystallographic groups [1]. Other than T_h , which swaps the two tetrahedra inside the cube of Figure 2, each is either the full symmetry group or a group of rigid motions of a Platonic solid. Dual Platonic solids share the same symmetry group; for instance, the icosahedron and dodecahedron share the same symmetries. Our list is complete if we name the chiral octahedral and icosahedral groups O and I , with the mirror versions called O_h and I_h .

Figure 5. A sphere painted with T_h symmetry and the mirror-symmetric source photo.

To keep this article a manageable length, we leave full consideration of these cases to the interested reader, mentioning only the detail that the 60-element icosahedral group I (the rigid motions) is generated by introducing into T an additional generator of order 5,

$$R_5 = \frac{1}{2} \begin{pmatrix} \phi & \phi^{-1} & 1 \\ \phi^{-1} & 1 & -\phi \\ -1 & \phi & \phi^{-1} \end{pmatrix},$$

where $\phi = (1 + \sqrt{5})/2$, the golden ratio. (It is a challenge to verify that R_5^5 is the identity.)

Experimenting with images invariant under the icosahedral and octahedral groups gives us visual evidence of the dual polyhedra. For instance, in images invariant under the octahedral group O , one can interpret the centers of 4-fold rotational symmetry as vertices of an octahedron or as face-centers of a cube. Similarly in I -invariant images, one sees 5-centers as either vertices of an icosahedron or centers of a dodecahedral face.

The technique for producing invariance under these groups is no different: one simply averages over the larger groups. These are programmed into the software, though icosahedral

symmetry—requiring averaging over a much larger group—sometimes causes a jerky motion when the user tries to rotate the sphere for a better view. We offer examples of each in Figure 6.

Accepted Manuscript

Figure 6. Screenshots showing octahedral symmetry, on the left, and icosahedral symmetry on the right. The octahedral painting draws on the sunset photograph from Figure 3. The icosahedral one uses a photograph of a single wisteria blossom, heavily post-processed to enhance the shapes.

Using source images with mirror symmetry also works to paint the sphere with functions invariant under the groups I_h and O_h , the full groups of symmetries of the icosahedron/dodecahedron and cube/octahedron.

5 COLOR-REVERSING SYMMETRY

Perhaps you wondered why we failed to mention odd functions in our introductory remarks about group averaging. An *odd* function on $\mathbb{R}^3$ is one that satisfies

$$f(-x) = -f(x),$$

where the evocative name reminds us that functions whose powers series consist entirely of odd powers of x . We can recognize odd functions by the symmetry of their graphs, which are invariant under a rotation of 180° about the origin.

Now that we have learned that a function value can be a color, it is not so great a stretch to think of the negative of that value as an opposite color, in some sense. There are various ways to implement this idea [2]. For the purposes of this paper, we do arithmetic modulo 256 on color values, so that, for instance, the negative of a color with red value R becomes $-R \equiv 256 - R \pmod{256}$. This can also be done with the Inverse function of Photoshop.

The trick we played with mirrors works just the same way to create colorings that have *color-reversing* symmetry. We say that a symmetry of the sphere, k , is a color-reversing symmetry of a function f when

$$f(k\vec{x}) = -f(\vec{x}), \text{ for all } \vec{x} \in S^2.$$

As an exercise, we invite you prove that

$$G = \{g \mid f(g\vec{x}) = f(\vec{x})\},$$

the symmetry group of the function f , is a normal subgroup of

$$G_c = \{h \mid f(h\vec{x}) = \pm f(\vec{x})\},$$

which we call the *color group* of f . The index of the symmetry group in the color group is always 2. A key ingredient of the proof is that the normalcy allows us to shuffle an actual symmetry past a color-reversing one to exchange it for a different color-reversing symmetry.

We show one example and simply mention the 6 other types, which the reader may enjoy investigating. Suppose we wish to create a coloring function whose symmetry group is T and

whose color group is O , the (chiral) symmetries of the cube/octahedron. We rig an image, on the left in Figure 7, where, once it is transferred to the sphere, a rotation through $\pi/2$ will reverse the colors. As you can see, this color-reversing feature again survives the group averaging. (On the advice of a referee, we experimented with reducing the area on the sphere covered by the source photograph.)

Figure 7. A polyhedral painting with color-reversing symmetry of type O/T , produced by a suitably chosen source image of butternut squash and its negative. The white and black stripes provide neutral regions to improve the lucidity of the final painting.

In general, we say an image has *type* G_c/G when its symmetry group is G and its color group of G_c , so the image in Figure 7 has type O/T .

What other types are there? One must go through each of the seven polyhedral groups we have studied and find out which can be normal subgroups of others. It turns out there are seven types: In Schoenflies notation, these are T_d/T , T_h/T , O/T , O_h/T_d , O_h/T_h , O_h/O , and I_h/I . With the seven symmetry types and seven more types of color-reversing symmetry, there is rich ground for artistic investigation.

6 IMPLEMENTATION AND CONCLUSION

We have highlighted the mathematical ideas behind polyhedral painting, but there are also interesting programming issues to solve, should faculty wish to assign it as a project for students. Tsao and Medal chose to use WebGL for their user interface, but there are many other choices.

Today's students hope to leave college having designed a software package with a graphical user interface (GUI); this project can be scaled to offer challenges anywhere from producing single images to a full-scale software product to paint polyhedra with all seven types of 3D point groups, using symmetry and color-reversing symmetry. Those pursuing the larger project are likely to discover different ways than ours to implement color-reversing symmetry.

The first design consideration is how to assign pixels from the source photograph to points on the sphere. The analytic geometry of this is not difficult, nor is the programming, especially with packages like WebGL or OpenGL that streamline many of the graphics tasks.

Next, the designer needs a way to manipulate colors as data types, typically using a vector (R, G, B) . Again, this can be difficult or easy depending on what software one uses for implementation.

An interesting sidelight of our implementation is an effort to boost saturation of averaged images. We adapted an earlier idea, [3], to define the function:

$$C(R) = \text{int} \left(128 \left(\frac{R-128}{128} \right)^{1/3} \right) + 128.$$

This drives Red color values greater than 128 higher and values lower than 128 lower. Replacing (R, G, B) with $(C(R), C(G), C(B))$ has the practical effect of increasing saturation, solving the problem of graying-out after averaging. For a future implementation, we propose replacing the power $1/3$ with a slider-controlled variable p . A p -value of 1 would cause no saturation adjustment; as p slides toward 0, the adjustment becomes more and more extreme.

The last programming ingredients one needs are matrix multiplication and management of loops to implement the various averaging routines. Depending on the degree to which coset averaging is used to reduce the number of terms in the average, this can be a brute force task or a more thoughtful one.

Tsao keeps a working version of the WebGL code at <http://rtsao.github.io/polyhedral-symmetry/>, along with the source code. We hope that readers will enjoy making their own polyhedral paintings as much as we have.

REFERENCES

- [1] Conway, J.H., H. Burgiel, and C. Goodman-Struass 2008. *The Symmetries of Things*. Wellesley, MA : A.K. Peters.
- [2] Farris, F. 2015. *Creating Symmetry: The Artful Mathematics of Wallpaper Patterns*. Princeton, NJ : Princeton University Press.
- [3] Farris, F. 2008. A Window on the Fifth Dimension. *MAA Focus* 2008 (Nov.): 1–4.
- [4] Grünbaum, B. and G.C. Shephard 1987. *Tilings and Patterns*. New York: W. H. Freeman.

BIOGRAPHICAL SKETCHES

Frank A. Farris served as Editor of *Mathematics Magazine* from 2001–2005, and again in 2008. He studied at Pomona College and then received the Ph.D. from M.I.T. in 1981. He has taught at Santa Clara University since 1984, after serving as Tamarkin Assistant Professor at Brown University. In Fall 2011, Farris visited Carleton College as Benedict Distinguished Visiting Professor. He has received the David E. Logothetti Teaching Award from SCU and the Trevor Evans Award from the MAA.

Ryan Tsao graduated *summa cum laude* from Santa Clara University in 2014, with majors in Computer Science and Computer Engineering, with minors in Mathematics and Art History. A native of Portland, OR, Tsao began his undergraduate studies at Rensselaer Polytechnic Institute before transferring to Santa Clara in 2011. He was elected to Phi Beta Kappa on graduation.