Article:
Hollerbach, R. (Accepted: 2025) Patterns on an Octahedron. Math Horizons. ISSN: 1072-4117 (In Press)

Patterns on an Octahedron

Rainer Hollerbach, School of Mathematics, University of Leeds

Most readers will be familiar with Rubik's cube, and may even be aware that similar puzzles exist with non-cubic shapes. The various manufacturers advertise these puzzles as being "educational for children six years and up." I suspect that might be a bit optimistic, but there is no doubt that for university and even high school students there is a great deal of interesting math associated with some of these puzzles and the patterns one can create on them.

The particular patterns we're interested in here are so-called color equality and equivalence patterns, which I had previously explored on Rubik's cubes (eprints.whiterose.ac.uk/id/eprint/199822, eprints.whiterose.ac.uk/id/eprint/211314, eprints.whiterose.ac.uk/id/eprint/229635). Color equality simply means that all six colors must be distributed exactly equally. That is, if you had two identically arranged cubes, you could pick any one color on the first cube, and any other color on the second cube, and you could then rotate the two cubes with respect to one another in such a way that the two colors coincide in all of their positions. If this holds true for *all* choices of colors on the two cubes, that arrangement satisfies the color equality criterion.

Figure 1(a) shows an example, on the simple $2 \times 2 \times 2$ cube. The left panel shows the three 'front' faces, then in the right panel the cube is turned over to show the three 'back' faces. We can easily see how red on the front matches up perfectly with orange on the back, similarly yellow matches with green, and blue with white. We also notice a three-fold rotational symmetry. That is, if you rotate the cube about an axis going through the front and back corners, you can match up red with either yellow or blue on the front, and orange with either green or white on the back. This pattern therefore satisfies the requirement that all six colors are distributed exactly equally.

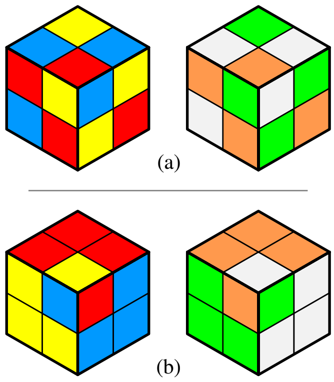

Figure 1: Two possible arrangements of a $2 \times 2 \times 2$ Rubik's cube: (a) an equality pattern, (b) an equivalence pattern.

Consider next the arrangement in Figure 1(b). Now red on the front does not match orange on the back, which is already enough to establish that this is not an equality pattern. The ways that red and orange are distributed are nevertheless very closely related: they are mirror images of each another. Furthermore, everything else we did above still applies as before. We still have the same three-fold rotational symmetry, so red on the front will still match either yellow or blue, and orange on the back will still match either green or white. That is, what we have here is a pattern where equality does not hold among all six colors, but if we split the colors into two sets, here taken as {red, yellow, blue} and {orange, green, white}, then: (a) within each set those three colors are distributed equally, (b) if you compare a color from the first set with a color from the second set, those two distributions are a mirror image pair. It is arrangements of this type that I called color equivalence patterns.

In the previous work I then explored all possible equality and equivalence patterns on $2 \times 2 \times 2$ and $3 \times 3 \times 3$ Rubik's cubes, studied their various symmetries, looked for interesting comparisons between different patterns, etc. What we want to do here is explore the same concepts on an octahedron. Figure 2 shows an octahedral puzzle called the Skewb Diamond, consisting of six corner pieces that have four colors each, together with eight central pieces having a single color.

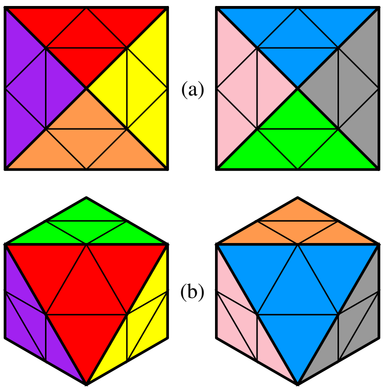

Figure 2: Four views of the Skewb Diamond puzzle. The (a) panels show 'top' and 'bottom' views along a corner-to-corner axis, as in Figure 1. The (b) panels then show the (a) positions tilted forward to lie flat on one face. Note the four-fold symmetry from the (a) perspective, and the three-fold symmetry from the (b) perspective. Cubes similarly have a three-fold symmetry when viewed as in Figure 1, and a four-fold symmetry when viewed straight down on one of the faces.

Now, as a puzzle the Skewb Diamond is somewhat limited; too many arrangements that one thinks ought to be accessible cannot actually be achieved by the allowed rotations. The mathematical analysis of which arrangements of various puzzles are achievable is itself a fascinating topic, but not what we want to consider here. Instead, let's just take this puzzle as our inspiration, but imagine that we are allowed to disassemble and reassemble it at will, so that all permutations and rotations of the six corners can be achieved, and similarly all permutations of the eight centers. What are then the possible equality and equivalence patterns, and what interesting symmetries and other properties might they have?

To be sure we really have *all* the patterns we want, it is necessary to write a computer program that scans through all possible arrangements and picks out the right ones. There is supplementary material at the end that outlines how to formulate such a code, but it is also possible to understand the patterns without considering how they were obtained. The only thing necessary to understand all the results to follow is a willingness to take the various templates also available as a separate document, cut them out and tape them together to form octahedra, and make sure you understand the various symmetries of the resulting three-dimensional patterns. There will also be a few challenges, which I urge you to spend some time trying to work out yourself before turning to the supplementary material for the solutions.

Equality Patterns

Including the original configuration in Figure 2, which itself constitutes pattern 1, there are 16 equality patterns. Figure 3 shows the next three. Patterns 2a,b constitute a mirror image pair. All six corners remain in their original positions, and rotate by either 90° (2a) or −90° (2b), where we will follow the usual convention that positive angles are measured in a counter-clockwise direction. For pattern 3 the corners again remain in their original positions, but now rotate by 180°. Since ±180° denote the same rotation, this pattern has no handedness, and is its own mirror image. From the way that these patterns are constructed it is clear also that they have no preferred axis; that is, the patterns look essentially the same when viewed along any of the axes.

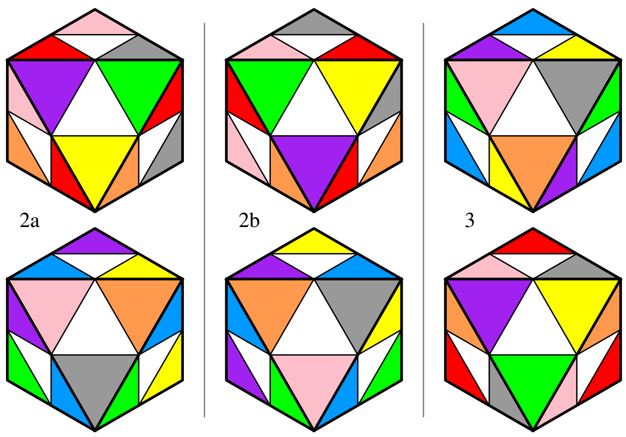

Figure 3: Patterns 2a,b and 3, as seen from the Figure 2(b) three-fold symmetry view. The orientations shown here make it particularly easy to see that red and blue are distributed equally, but the other six colors are also the same. Also, we are considering only the six corners so far, so the centers are left blank.

Figure 4 shows the next two patterns, and compares them also with pattern 3. Both patterns 4 and 5 yet again leave all six corners in their original positions; pattern 4 then rotates the top and bottom corners by 180°, whereas pattern 5 rotates the other four corners by 180°. That is, we could also think of pattern 3 as being a superposition of patterns 4 and 5, which is indeed reflected in

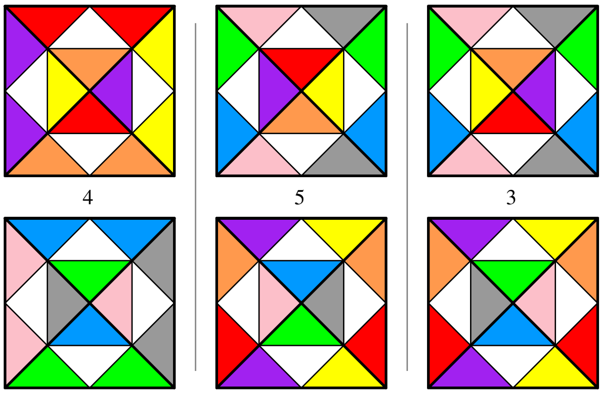

Figure 4: Patterns 4 and 5, also compared with pattern 3, all seen from the Figure 2(a) four-fold symmetry view.

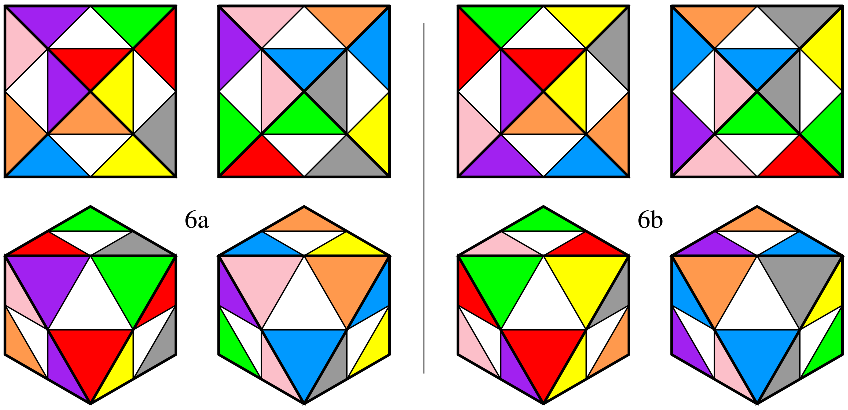

Figure 5: Patterns 6a,b, seen from both the Figure 2(a,b) views.

the figure. Unlike pattern 3 though, which has no preferred axis, patterns 4 and 5 clearly do have a preferred axis, namely which two opposite corners you choose to constitute the top and bottom. Really there would therefore be three possibilities for these patterns, depending on which axis we choose. These are all basically the same patterns though, so we will consistently only choose one of the three axes.

Figure 5 shows another mirror image pair, pattern 6a,b. This pattern has both a preferred axis and handedness, that is, it forms a mirror image pair. In this case it is left up to you to print out the provided templates and work out yourself how the pieces move to create these patterns.

Overall there are then four possible symmetries: either a preferred axis or no preferred axis, and either handedness or no handedness. In addition to the eight patterns shown here, there are four other mirror image pairs, with all four of them also having a preferred axis. By studying the patterns given here, can you perhaps guess what these other four mirror image pairs are? The complete set of all 16 patterns follows remarkably simple rules, but still yields quite a variety of patterns, including all four of these symmetry types.

Equivalence Patterns

Based on the results so far, and comparing also with Figure 1(b), can you perhaps already guess at least a few equivalence patterns? The full set is presented in the supplementary material. Before we even get to any of the details though, equivalence patterns raise a question that fundamentally does not arise for equality patterns: how do you divide the colors up into the two mirror image sets?

Imagine a cube, without any colors attached to the faces at all yet. In how many different ways can you divide the faces into two sets of three? From a combinatorics point of view, the answer would be: choose three out of six, so $6!/3!^2 = 20$. However, most of these are just rotated versions of one another. From a geometric point of view, there are only two different ways of dividing the faces up into two sets of three.

First, as shown in Figure 6(a), you can split every pair of opposite faces between the two sets. Alternatively, as in Figure 6(b), you can split only one pair of opposite faces between the two sets, and have the other two pairs of opposite faces each entirely within a given set. Again, any of the 20 combinatorics choices of three of the six faces on a cube will always be just some suitably rotated version of one of these two ways of splitting things.

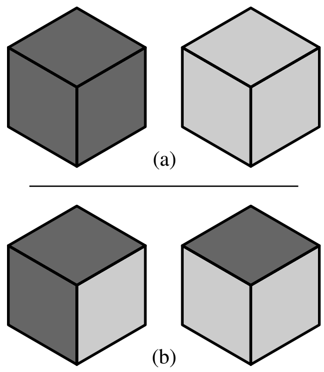

Figure 6: The two different ways of dividing the six faces of a cube into two sets of three faces each.

Returning to Figure 1(b), it clearly divides the six colors as in Figure 6(a). In fact, all equivalence patterns on a cube split the colors up in this way; there are no equivalence patterns where the colors are split up as in Figure 6(b). Why should that be the case? At one level we could just say: because the computer code checked every possibility, and there simply aren't any patterns that follow the 6(b) way of dividing things. At a deeper level though, we could at least conjecture that it is due to the fact that the 6(a) split is more symmetric than the 6(b) split, with every pair of opposite faces being treated equally, rather than one pair being treated differently from the other two.

So, this raises two obvious points: First, how many different ways are there of dividing the eight faces on an octahedron into two sets of four each? That is, what is the equivalent of Figure 6 for an octahedron? See if you can work out the possibilities here. There are definitely more than just two!

Next, once you think you have them all, apply symmetry considerations as above to try to decide which splitting options are most likely to yield equivalence patterns. The results are presented in the supplementary material, but these fundamental questions of what is the equivalent of Figure 6 for an octahedron, and which options are most likely to yield equivalence patterns, are something you should definitely think about yourself first.

Adding the Centers

We are rapidly running out of space, and haven't considered the centers at all yet, so let's keep this section short. Figure 7 shows two possible arrangements of the centers, together with the corners in their pattern 1 original configuration. Both arrangements yield color equality patterns. If you focus on the distribution of just a single color at a time, the two arrangements even look the same: in both cases each color has its center moved to one of the immediately adjacent faces. Nevertheless, these arrangements are different. This becomes especially clear if we think of them as actions being applied to the centers, and ask how often the actions would have to be repeated before everything returns to its original configuration. In Figure 7(a) you have to repeat the action four times to get back to where you started, whereas in Figure 7(b) centers are exchanged pairwise, so repeating that action twice gets you back where you started. Next, apply these same center arrangements to some of the other equality patterns. Are the results still equality patterns? If not, why does it work for pattern 1 but not for some of the others?

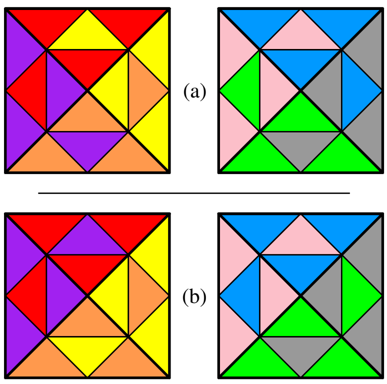

Figure 7: Two different arrangements of the centers, with the corners in their original positions.

Future Work

You might also enjoy Roice Nelson's "Abstracting the Rubik's Cube" (*Math Horizons*, April 2018), where he discusses how various puzzles such as these can be extended into higher dimensions, at least as mathematical abstractions. The concepts of color equality and equivalence patterns would still apply, but actually computing these patterns for even just the simplest such extension, the four-dimensional $2 \times 2 \times 2 \times 2$ hyper-cube, would be far more challenging than anything we considered here.

Supplementary Materials

As indicated in the main article, there are several items of supplementary material. First, there is a separate document containing templates to cut out and fold together to create all of the cubes and octahedra in the main article. All figures, both in the main article and in the later sections here, were chosen to show the given patterns as well as any two-dimensional representation can, but we are ultimately dealing with three-dimensional patterns here, spread out over the surfaces of objects having six (cubes) or even eight (octahedra) sides. It is inevitable then that at least some of the resulting patterns will be sufficiently intricate that you really do have to tape together the given templates to see for yourself what the actual three-dimensional patterns look like when viewed from whatever angles you want.

In some cases it might even be helpful to print out two copies of a given pattern, to compare one color as seen from one orientation with another color as seen from a different orientation. We recall after all the original definition of color equality: If you had two identically arranged objects, you could pick any one color on the first, and any other color on the second, and rotate the two objects with respect to one another so that the two colors coincide in all their positions. So if you're not quite seeing the claimed results on just a single copy alone, then explicitly do exactly what this original definition says to do, and compare two copies viewed from different orientations.

The article then also presented a number of challenges:

1. Can you guess the eight equality patterns that weren't given in the article, and thereby obtain the – surprisingly simple – rules governing all 16 of them?

2. Can you guess at least a few of the equivalence patterns? The rules here turn out to be considerably more complicated than for the equality patterns, so it is very unlikely that you would be able to guess all of the equivalence patterns, but some of them are sufficiently simple – and similar to some of the equality patterns – that you might be able to guess those.

3. As the first step in systematically classifying all possible equivalence patterns, what is the equivalent of Figure 6 on an octahedron? How many geometrically different ways are there of splitting the eight faces of an octahedron into two groups of four each? And once you think you have them all, which ones are like Figure 6(a) in treating every pair of opposite faces the same, and which ones are like Figure 6(b) in treating some pairs of opposite faces differently from others? How many fundamentally different types of equivalence patterns might we therefore expect?

4. Given at least some of the possible corner arrangements, both the ones in the article as well as any others you may have guessed, play around a bit with different arrangements of the centers, both as given in Figure 7 as well as other arrangements you can make up yourself. If a given center arrangement works for the original corner pattern 1, will it also work for other patterns? Or if not, why not? The answers will invariably have something to do with different corner patterns having different symmetries. Forcing you to think about the symmetries of these various patterns is precisely what makes this entire problem so appealing though.

Note also that to really work through all of these challenges yourself is not trivial, and is likely to take a considerable amount of time. I would strongly encourage you though, do take the time. Print out enough of the templates that you really do understand the given patterns, and then spend some time thinking about these four challenges here. The best way to learn anything related to math is always to try to work things out for yourself first before turning to solutions provided by someone else. I hope also that you might find some of these challenges genuinely enjoyable.

The remainder of these supplementary materials then consists of four sections on: (a) equality patterns, (b) equivalence patterns, (c) center arrangements, as applied to both equality and equivalence patterns, and (d) an outline of how the computer scan of all possible arrangements works. There are also further templates, again in the separate template document.

You will find that section (b) on the equivalence patterns involves some rather intricate details and symmetries, which will probably take some time to fully grasp. It is precisely by exploring all the various cases though that you will learn something about some of the possibilities that can arise with symmetries of three-dimensional patterns. And of course, I hope you will find it interesting to see how some of these symmetries come together to create fundamentally different types of equivalence patterns.

Section (d) on the computing aspects also involves some technical details that may not necessarily appeal to all readers. However, anyone not interested in computing can skip this section without missing out on any of the rest of the material. For any readers who are particularly intrigued by computing though, this outline should provide enough details to allow you to develop your own codes if you wanted.

Equality Patterns

As colorful as they undoubtedly are, the full set of 16 equality patterns follows remarkably simple rules. First, choose one of the three pairs of opposite corners to constitute the 'top/bottom' pair. As in the article, we take the red/purple/orange/yellow corner as the top, and the blue/pink/green/gray corner as the bottom. These two corners will follow one rule, whereas the other four corners will follow another rule. These two rules are: rotate the top/bottom pair together, and similarly rotate the other four together. That is, let α be the rotation angle of the top/bottom pair, with the allowed values being $0°$, $90°$, $-90°$, $180°$. Similarly let β be the rotation angle of the other four corners, with the same allowed values.

β \ α	$0°$	$90°$	$-90°$	$180°$
$0°$	1	7a	7b	4
$90°$	6a	2a	10a	8a
$-90°$	6b	10b	2b	8b
$180°$	5	9a	9b	3

Table 1: The 16 equality patterns with their associated $\{\alpha, \beta\}$ values.

Table 1 then lists the 16 possible combinations, each labelled according to the pattern numbers that we are associating with those particular $\{\alpha, \beta\}$ values. If we compare in particular the values for patterns 1, 2a,b, 3, 4 and 5 with the previous descriptions in the article, everything exactly matches. Patterns 6a,b, where you were asked to work out yourself how the pieces move around, are also seen to correspond to $\alpha = 0°$, $\beta = \pm 90°$. The four missing pairs 7, 8, 9, 10 then fill out the table.

Figure 8 shows the complete set, where we note also that it is easiest to separately compare all the top halves on the left, and then the bottom halves on the right. We clearly see how the top/bottom pieces vary from column to column (α), whereas the other four pieces vary from row to row (β). Once again, you are also encouraged to print out at least a few of the corresponding templates, to convince yourself that all of these really are equality patterns. In some cases it is already clear in the figure, but for patterns like 6a,b or 8a,b, for example, the eight colors are sufficiently mixed up that it is probably not obvious just from this figure that these patterns satisfy the equality criterion.

Once we recognise these simple rules for the full set, all of the previously noted symmetries also fall into place. First, if $\alpha = \beta$, so patterns 1, 2a,b, 3, then all six pieces are actually being treated in exactly the same way, so there is no preferred axis after all. In contrast, if $\alpha \neq \beta$, then the top/bottom pair is different from the other four, and thus constitutes a preferred axis. Next, if we allow only the values $0°$ and $180°$ for both angles, those patterns (1, 3, 4, 5) will be their own mirror images. In contrast, any pattern that has a value of $\pm 90°$ for either of the angles will

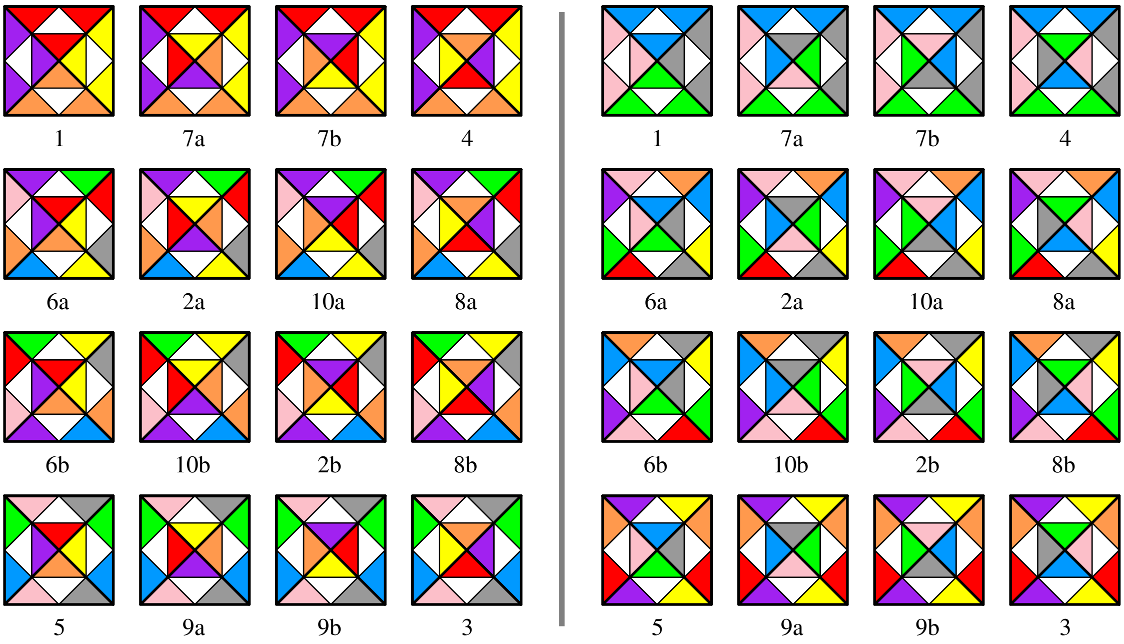

Figure 8: The 16 equality patterns as in Table 1, with top/bottom halves on the left/right.

necessarily occur as a mirror image pair, so 2a,b, 6a,b, etc. By using some of the templates, check also that these really are mirror image pairs in all cases. That is, if you compare, for example, red on 8a and 8b, those two distributions really are mirror images of each other.

There is one further symmetry we can note in Figure 8: Every one of these patterns has a four-fold rotational symmetry. That is, if you rotate the entire patterns by 90° for the top views, and correspondingly by −90° for the bottom views, then the colors consistently map into each other as red → purple, blue → gray, etc. This symmetry comes about because the 'β-rule' states that those four corners all rotate in exactly the same way.

In principle one could also imagine having three separate rules whereby all three pairs of opposite corners behave differently, in which case we would expect to lose this four-fold rotational symmetry. For the equality patterns it turns out that there just aren't any patterns having three separate rules like this, but we will see below that there is one equivalence pattern where we do have three separate rules, and the four-fold rotational symmetry is then indeed lost, and replaced by just a two-fold symmetry.

Finally, we could ask, if we didn't specify this particular choice of top/bottom pair, what would then be the total number of equality patterns? It is straightforward to work out the answer here. All patterns except 1, 2a,b and 3 have a preferred axis, so for all those patterns there would be three possible choices. For patterns 1, 2a,b and 3 though there is only this one choice, since these would look the same even if you started out with a different choice of top/bottom pair. The total number of equality patterns is therefore 4 (coming from 1, 2a,b, 3) plus 3 times 12 (coming from all the others), for a total of 40 equality patterns, if we want to include all these variants that are technically different even though they are effectively the same patterns, just with different top/bottom choices.

Equivalence Patterns

There are two preliminary points regarding equivalence patterns, on either a cube or an octahedron. First, we saw before that equality patterns could either lack handedness (1, 3, 4, 5) or occur as mirror image pairs (all the others). Is the same also true for equivalence patterns? Well, what would a non-handed equivalence pattern look like? If you think about it a bit, you realise that a non-handed equivalence pattern is just the same as a non-handed equality pattern. That is, take something like pattern 4, for example, and divide the colors up into the two sets {red, purple, orange, yellow} and {blue, pink, green, gray}. Since we already have equality among all eight colors, we trivially also have equality within each of these sets separately. Because pattern 4 lacks handedness though, we could also describe the relationship between red and blue, say, not as 'equality' but rather as 'mirror images'. After all, 'lacks handedness' means precisely 'is it's own mirror image'.

So indeed, any non-handed equality pattern can automatically also be thought of as a non-handed equivalence pattern. Now, in some contexts it can actually be important to include these patterns. For example, for the group theoretical analysis of the Rubik's cubes equivalence patterns (eprints.whiterose.ac.uk/id/eprint/229635), it was necessary to include these non-handed patterns, in order to have a well-defined group that contains all the required elements. However, when considering patterns from a purely geometrical perspective, as we're doing here, it would be rather pedantic to insist that patterns 1, 3, 4 and 5 can also be thought of as equivalence patterns. For present purposes, therefore, we simply assert that these non-handed patterns do not count as equivalence patterns, and we are only interested in distinct mirror image pairs.

The second point is specifically about the nature of mirror image pairs, and how they differ between equality and equivalence patterns. Let's return to Figure 1 from the article, showing simple Rubik's cube examples of equality and equivalence patterns. For convenience, these are shown again in the left half of Figure 9. The right half shows their respective mirror image partners. Now, let's

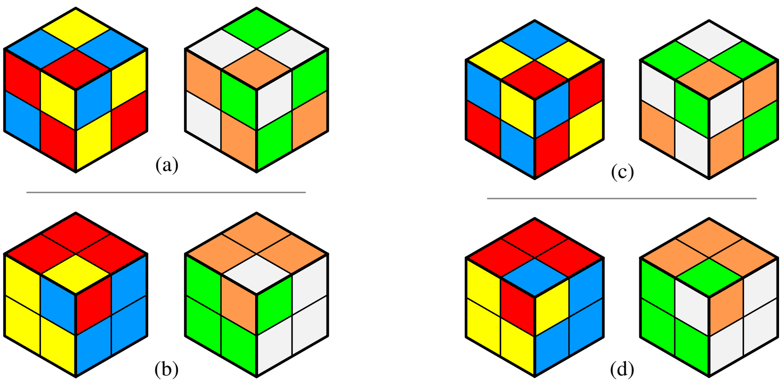

Figure 9: The Figure 1 patterns shown again on the left, together with their mirror image partners on the right.

compare the two equivalence patterns, panels (b) and (d). On the one hand, we very clearly recognise that they are indeed a mirror image pair. On the other hand, we can also match them up directly, by simply comparing the red/yellow/blue 'front' of (b) with the orange/green/white 'back' of (d), and vice versa. That is, unlike the equality mirror image pairs (a) and (c), which cannot be matched up in any way, equivalence mirror image pairs can be matched to each other, if you just suitably interchange the 'front' and 'back' colors.

In some sense therefore we can think of (b) and (d) as being the same pattern after all, in which case we could dispense with the whole idea of mirror image pairs entirely, and only count this as one pattern. This is very similar to the concept of preferred axes that we saw before – and here as well, by the way – where we had effectively the same patterns three times. Unlike the preferred axes though, where we chose to include the patterns only once rather than three times, for the mirror image pairs we will explicitly include both versions in all the equivalence patterns to come.

The six divisions, leading to Types I, II, III

With these preliminary items dealt with, we next turn to your main challenge, namely to construct the equivalent of Figure 6 for an octahedron. As shown in Figure 10, there are six different ways of dividing the eight faces of an octahedron into two sets of four each. The first two, (a) and (b), follow directly from our very first view of the octahedron, in Figure 2. For either the (a) or (b) perspectives in Figure 2, simply assign the four top faces to one set, and the four bottom faces to the other set, thereby yielding the (a) and (b) divisions here.

For the next division, (c), you should definitely print out the corresponding template, since it is not at all obvious from this representation here what it actually looks like when spread out over the surface of the octahedron. Once you've constructed the three-dimensional pattern though, you can easily see that it is basically the equivalent of the Figure 6(b) division of the cube. In both cases the cube or octahedron gets divided up into two sets very similar to the two pieces that constitute the surface of a standard American baseball.

Division (d) divides the octahedron up into four equal wedges, like the slices of an orange. Division (e) divides it up into eight separate octants, in such a way that no two adjacent faces are assigned to the same set. Finally, division (f) is a rather ugly, lopsided way of splitting things up.

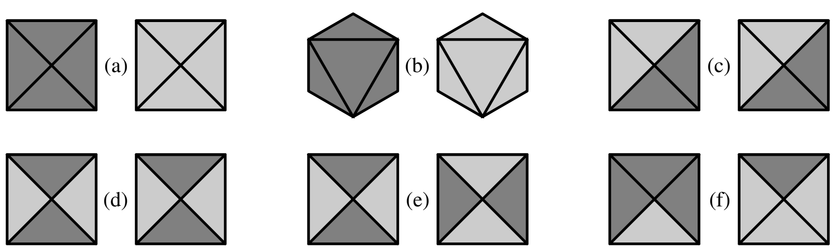

Figure 10: The six possible ways of splitting the faces of an octahedron into two sets of four each.

And again, these six divisions really are the only ways of dividing the eight faces up into two sets of four each. Any other choice of four out of the eight faces will always be just some suitably rotated version of one of these six divisions.

So, which of these divisions are sufficiently symmetrical that they are likely to have associated equivalence patterns? Well, if the Figure 6(b) division didn't work for the cube, then for exactly the same reason the (c) division here probably won't work either. As elegant as they both may be in some regards, they both have the property that some pairs of opposite faces are treated differently from other pairs, with some being split between the two sets but others not. The ugly division (f) can be excluded for the same reason.

Let's consider next the division (b), which is the one that is most borderline in terms of whether it's likely to work or not. In particular, every pair of opposite faces is treated the same, being split between the two sets. However, they're still not all on an entirely equal footing. The one pair of opposite faces that constitutes the top/bottom in division (b) has a slightly different status from the other three pairs, since its top part is completely surrounded by three others that belong to its same set, and similarly its bottom part is also surrounded by three others belonging to its set. None of the other three pairs of opposite faces has this property. So, is just this slight asymmetry already enough to exclude the possibility of obtaining equivalence patterns that follow this division? Well, the computer scan of all possible arrangements didn't reveal any 'Type (b)' equivalence patterns, so perhaps this slight asymmetry is indeed the reason?

The remaining three divisions (a,d,e) are all completely symmetrical in every regard, with nothing that would distinguish any pair of opposite faces from any of the others. We suspect therefore that all three of these ought to yield equivalence patterns, which does indeed turn out to be the case. So, first, let's relabel (a,d,e) as Types I, II, III. There is then one geometrical difference we can note between Types I and II versus Type III. Figure 11 shows what these three divisions look like, first from the same 'top/bottom' perspective as in Figure 10, labelled the (1) view in Figure 11. Next, the (2) and (3) views in Figure 11 show what these divisions look like when you orient them differently, to view them along the other two possible axes. Here again it is probably helpful to print out the templates, construct the actual three-dimensional octahedra, and rotate

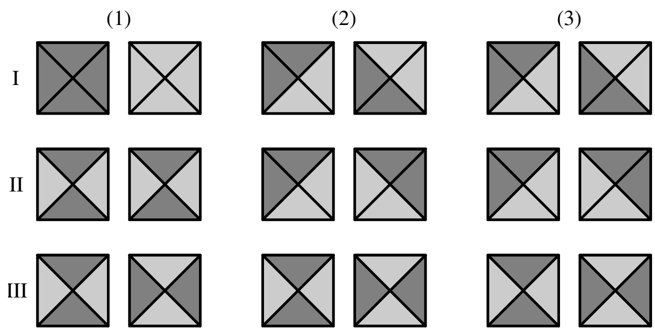

Figure 11: The three divisions Types I, II, III, seen from all three axes.

them into the appropriate orientations to more directly see what Figure 11 is showing.

Now, what is the significance of Figure 11? Types I and II look different from the (1) perspective versus the (2) or (3) perspectives, whereas Type III looks the same from all three. That is, Types I and II have a preferred axis, whereas Type III does not. This will be reflected in the equivalence patterns below, where Type I and II patterns will necessarily 'inherit' that preferred axis from their fundamental way of dividing the faces into the two sets of four each, whereas Type III patterns could exist that do not have a preferred axis.

	Set 1 Colors	Set 2 Colors
Type I	red, purple, orange, yellow	blue, pink, green, gray
Type II	red, orange, blue, green	purple, yellow, pink, gray
Type III	red, orange, pink, gray	purple, yellow, blue, green

Table 2: If you apply the dark/light divisions from Figure 11 to our original color scheme from Figure 2, you end up with the colors divided as indicated.

Finally, Table 2 explicitly indicates how the colors are divided into the two sets then, based on these geometrical descriptions of Types I, II and III, as applied to our octahedral color scheme from Figure 2. Note also that for Types I and II this assumes our previous red/purple/orange/yellow corner as the top, and blue/pink/green/gray corner as the bottom. If we had chosen a different preferred axis, the separation into the two sets would also be correspondingly different. For Type III though, since it has no preferred axis, the division into the two sets would always be as in Table 2, regardless of what you might choose for your top/bottom corner pair.

Type I Equivalence Patterns

Figure 12 shows patterns 11a,b and 12a,b, which are very closely related to the equality patterns 7a,b ($\alpha = \pm 90°$, $\beta = 0°$) and 9a,b ($\alpha = \pm 90°$, $\beta = 180°$), respectively. The only difference is that the top/bottom corners, which previously rotated in the same directions (that is, $\alpha = \pm 90°$ applies to both), now rotate in opposite directions. This is already enough to switch the patterns from equality to Type I equivalence. This is especially clear for patterns 11a,b, where the colors remain nicely separated between {red, purple, orange, yellow} on the top and {blue, pink, green, gray} on the bottom. And sure enough, each set separately is easily seen to satisfy equality, simply by applying the four-fold rotational symmetry that also still applies. If we compare, for example, red versus blue though, corresponding to top versus bottom, those distributions are indeed mirror images, as required for a Type I equivalence pattern. That patterns 12a,b follow the same symmetries is less obvious, since now the top/bottom colors are more mixed up, but if you print out templates and construct the actual octahedra, you can verify that 12a,b are indeed also Type I equivalence patterns.

Figure 13 shows the case alluded to above, where the three sets of opposite corners follow three

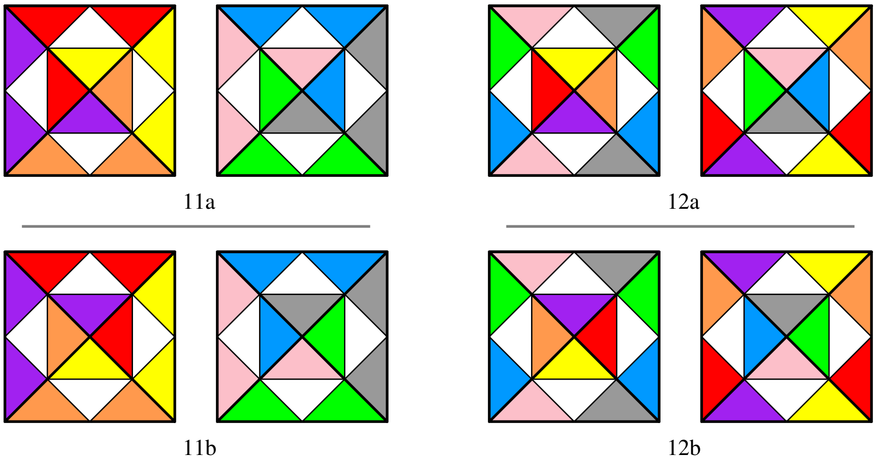

Figure 12: Patterns 11a,b and 12a,b.

different rules. In particular, the rules here are: (i) the top/bottom corners rotate $\pm 90°$ in opposite directions, just as they do for patterns 11a,b and 12a,b; (ii) of the two other pairs of opposite corners, one pair rotates $90°$ and the other pair rotates $-90°$. With these rules, there are now also four such patterns rather than just the usual mirror image pair. That is, rule (i) already gives you two choices, and rule (ii) independently gives you another two choices. To confirm that these patterns 13(a,b,c,d) are Type I equivalence patterns, with the colors divided as in Table 2, will again probably require the templates. For each version (a,b,c,d) half the colors are strictly separated on either the top or the bottom, so for those it is easy to see already which ones are equal and which ones are mirror images, but for the other half of the colors that are spread out across both top and bottom it is probably necessary to see the actual three-dimensional patterns. Note finally that the previous four-fold rotational symmetry is indeed lacking here, and there is only a two-fold rotational symmetry. Given rule (ii) this is exactly what we would expect; you have to rotate the whole pattern by $180°$ rather than just $90°$ for rule (ii) to still look the same.

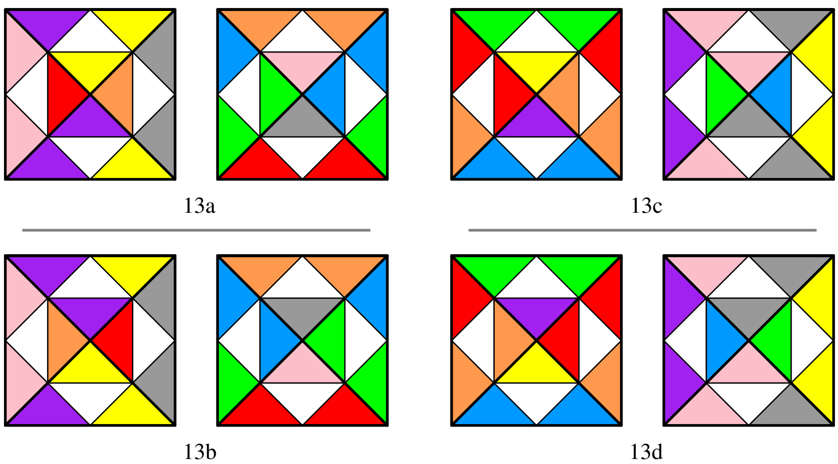

Figure 13: Patterns 13(a,b,c,d).

Type II Equivalence Patterns

All patterns so far, from the simplest equality patterns up to the rather complicated quadruple set 13(a,b,c,d), have left all six corners in place, and merely rotated them in various ways. It is remarkable that such a variety of patterns and symmetries can be created from actions as simple as this. Nevertheless, it would also be nice to have at least a few patterns that involve corners moving to new positions. Fortunately, all of the remaining patterns do exactly that. Figure 14 shows the patterns 14a,b and 15a,b, which are again very closely related to the previous equality patterns 7a,b and 9a,b, respectively. In particular, if we compare the four corners other than the top/bottom pair, they are identical between 14a,b and 7a,b, and between 15a,b and 9a,b. Next compare the top/bottom corners. Those are also identical, except they have exchanged positions, with top/bottom becoming bottom/top. Otherwise though, even the orientations are identical between top/bottom of 14a,b and bottom/top of 7a,b, and similarly for 15a,b compared with 9a,b.

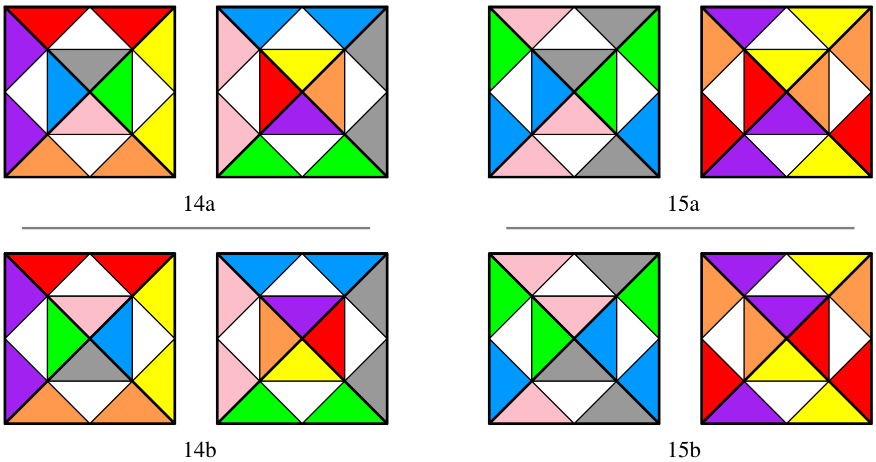

Figure 14: Patterns 14a,b and 15a,b.

So what is the effect of this top ↔ bottom corner interchange? It creates Type II equivalence patterns. This is most easily seen for patterns 15a,b, where the colors are again nicely separated between top and bottom. We can then clearly see that the two sets {red, orange, blue, green} and {purple, yellow, pink, gray} each separately satisfy equality, and if you compare for example red with purple, those two distributions are mirror images of one another. Comparing with Table 2, this is exactly the color division that defines Type II. That patterns 14a,b are also Type II is less obvious, since this is another case where the colors are more spread out over both the top and the bottom. Once again though, if you use the templates to construct actual octahedra, by suitably rotating them you can convince yourself that 14a,b also satisfy the Type II color division from Table 2.

As always, the most interesting aspects of these various patterns are the associated symmetries. There is indeed something here that we have not encountered before. These patterns lack the four-fold rotational symmetry that we had for all patterns thus far, with the exception of the slightly anomalous quadruple set 13(a,b,c,d). This lack of four-fold symmetry is most easily seen for patterns 15a,b. In particular, notice how all of the 15a,b panels have one diagonal where two same-color patches are next to each other, whereas along the other diagonal the colors are more separated. Well, if the two diagonals are different, then you would have to rotate the whole pattern by 180° to get the same arrangement back again. Rotating by only 90° would effectively interchange the 15a,b diagonal types. That is, rotating by 90° makes either pattern look like its mirror image partner.

Now, what causes this? Why do we lose the four-fold rotational symmetry, and instead only have a two-fold rotational symmetry? For patterns 13(a,b,c,d) the explanation was very simple: the corners that define the two diagonals followed different rules. Here though that cannot be the explanation. We recall in particular that 15a,b has its four non-top/bottom corners exactly the same as 9a,b, and for all 16 equality patterns there was only that one single 'β-rule' governing all four of these corners.

The explanation instead is that this is simply an inherent property of being a Type II equivalence pattern. In particular, return to Figure 11, and look at Type II from the (1) perspective, corresponding to viewing it along its top/bottom preferred axis. If you take that dark/light arrangement and rotate it by 90°, the dark and light wedges exchange places. What 'dark' and 'light' represent though are precisely the two mirror image color sets, {red, orange, blue, green} and {purple, yellow, pink, gray}. That is, taking *any* Type II equivalence pattern and rotating it by 90° about its preferred axis will *necessarily* make it look like its mirror image partner, which is exactly what we see with 15a,b. (This is also true for 14a,b, but less obviously so.)

Note incidentally that Type I equivalence patterns are quite different in this regard. If you look at Type I from the (1) perspective in Figure 11, and rotate that arrangement by 90°, the dark and light halves remain unchanged. And sure enough, patterns 11a,b and 12a,b have a four-fold rotational symmetry, and map into themselves rather than their mirror image partners if you rotate them by 90°. (And again, the reason why the quadruple patterns 13 lack the four-fold symmetry is specific to them, and not any inherent property of Type I equivalence patterns.)

Type III Equivalence Patterns

Patterns 16a,b, shown in Figure 15, are very closely related to patterns 14a,b. Rotating the top corners of 14a,b by 180° yields 16a,b, or equivalently rotating the bottom corners of 14a,b by 180° yields 16b,a. These rotations of one single piece are enough to turn the patterns from Type I to Type III equivalence. Figure 15 also shows two further sets, 16a',b' and 16a",b". These are the same patterns, but with different choices of preferred axis. We recall that any time there is a preferred axis, there will inevitably be three possibilities, but which are ultimately all the same pattern. For the previous patterns we only presented one choice, consistently taking red/purple/orange/yellow and blue/pink/green/gray as the top and bottom corners.

So why present all three choices here? Well, it certainly doesn't hurt to have one example anyway where we explicitly show all three choices. In particular, convince yourself that 16a',b'

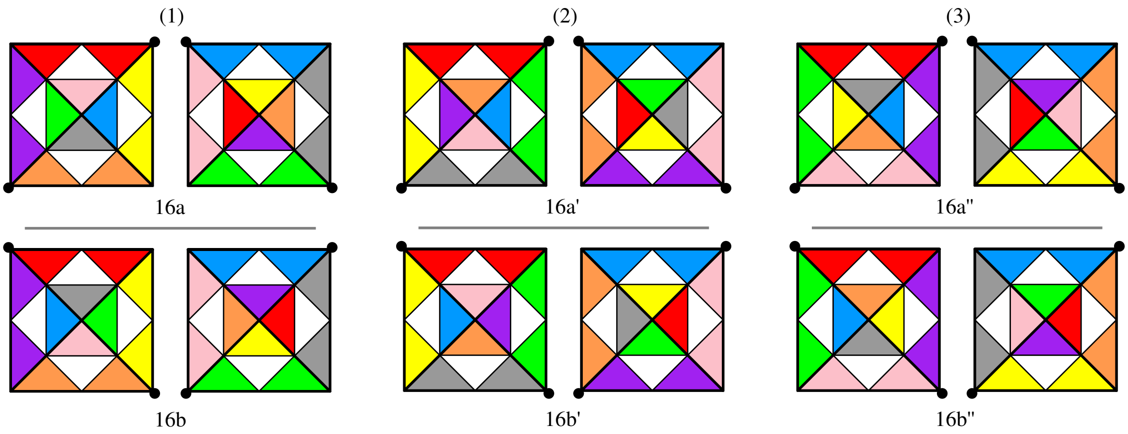

Figure 15: Patterns 16(a,b,a',b',a",b"). The numbers (1,2,3) at the top refer to the three possible choices of preferred axis. The black dots beside some of the corners indicate the ones that will be rotated by 180° to produce Figure 16.

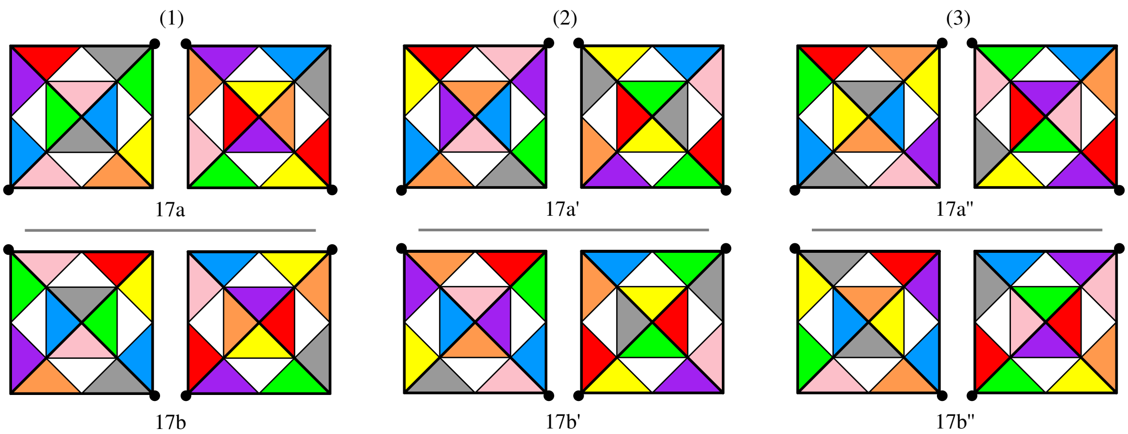

Figure 16: Patterns 17(a,b,a',b',a",b"), which end up satisfying 17a = 17a' = 17a" and 17b = 17b' = 17b".

and 16a",b" really are the same patterns as 16a,b, with the only difference being the choice of top/bottom corners. More importantly though, including all three choices here is an important intermediate step in what we want to do next. Specifically, for all six patterns shown in Figure 15, take the pieces indicated by the black dots beside them, and rotate those by 180°. The results are shown in Figure 16, and constitute patterns 17(a,b,a',b',a",b"). Make sure you compare patterns 16 and 17 sufficiently carefully that you really do understand how all six of patterns 17 are related to the corresponding patterns 16.

Now, in terms of how they were created, we would be tempted to say that patterns 17 ought to be like patterns 13, which we recall had two preferred axes. That is, for patterns 13 all three pairs of opposite corners followed different rules, which certainly also seems to be the case here: first the top/bottom pair are interchanged (and suitably rotated), next one of the other pairs is left in place but rotated by 180°, and finally the third pair is left alone entirely. So how could that be anything other than 'two preferred axes'? Except, it's not! There is *no* preferred axis here! Specifically, the three patterns 17(a,a',a") are actually all the same, and similarly 17(b,b',b") are all the same. Use the templates to convince yourself of this. Take the one single octahedron that has labels 17a, 17a' and 17a" on three of its faces, and sequentially orient it so it has those three labels on top. In each case it will agree perfectly with the 17(a,a',a") results shown in Figure 16. That is, while 16(a,a',a") really are different, and require three different templates, 17(a,a',a") are all the same, and require only one template.

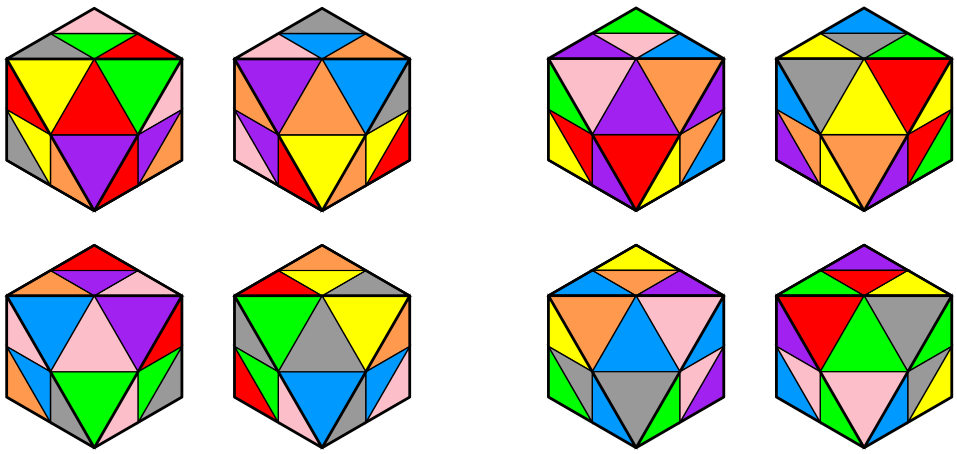

Figure 17: Pattern 17a from the three-fold symmetry view. The four panels on the left have {red, orange, pink, gray} on top, whereas the four panels on the right have {purple, yellow, blue, green} on top, in accordance with the Type III color division in Table 2.

This special property that patterns 17a,b have no preferred axis, even though the way they were constructed makes it seem like they should have two preferred axes, certainly makes them my favorites. And since they are so nice, let's also look at them from a different perspective, that best illustrates their status as Type III equivalence patterns. Figure 17 shows 17a from eight different orientations, so all eight colors in turn are on top, in the Figure 2(b) three-fold symmetry perspective. The centers have also been added here, in a way designed to indicate which color to focus attention on in each panel. Specifically, in each panel, focus attention on the color of the center on top, and look at where its three color patches coming from the corners are. In each case they form a 'swirl' around that center, with half the swirls going to the left and the other half to the right. That is, we separately have equality within the two sets {red, orange, pink, gray} and {purple, yellow, blue, green}, and if we compare red and purple, say, those two distributions are mirror images of each other. Comparing with Table 2, this is indeed exactly the color division that defines Type III. There is also a template for 17b with the centers included in the same way, which you can use to verify that the colors are divided in the same way, but now for each color the swirls go in the opposite direction as for 17a.

A final comparison

Let's start by comparing one of our earliest results, patterns 2a,b from Figure 3, with this latest result, pattern 17a from Figure 17. We notice that 2a,b exhibit exactly the same swirls as in Figure 17. The 'only' difference is that for either 2a or 2b all the swirls go in the same direction, whereas for 17a,b left-handed and right-handed swirls are combined in the same patterns. This is of course exactly what constitutes the difference between equality versus equivalence patterns. So, let's just summarise the complete list of patterns where the single-color distributions are the same mirror image pairs, but the patterns are nevertheless different:

- 2a,b (Equality), 17a,b (Type III Equivalence)
- 7a,b (Equality), 11a,b (Type I Equivalence)
- 9a,b (Equality), 12a,b (Type I Equivalence)
- 10a,b (Equality), 13a,b,c,d (Type I), 14a,b, 15a,b (both Type II), 16a,b (Type III)

In terms of single-color distributions, the 10a,b type is thus the most interesting, yielding the greatest number of different patterns, including even all three different types of equivalence patterns.

Adding the Centers

The article suggested taking the center arrangements in Figure 7 and applying them to some of the other patterns. Table 3 shows the results for the equality patterns 1-10. For 1-3 the results are still equality patterns, for 4 and 5 they are Type I and II equivalence patterns, respectively, for the Figure 7(a,b) arrangements, and for 6-10 these combinations of corners and centers satisfy neither equality nor equivalence. We can understand all of these results as follows:

We begin by noting that the center arrangements in Figure 7 themselves have a certain handedness associated with each color, in the sense that its center moves either clockwise or counterclockwise from its original position. For Figure 7(a) the top colors {red, purple, orange, yellow} all

	1	2a,b	3	4	5	6a,b	7a,b	8a,b	9a,b	10a,b
Fig. 7(a)	E	E	E	I	I	–	–	–	–	–
Fig. 7(b)	E	E	E	II	II	–	–	–	–	–

Table 3: The results of combining the center arrangements from Figure 7(a,b) as indicated with the ten equality patterns 1-10. E denotes an equality pattern, I and II are Type I and II equivalence patterns, and '–' means that those combinations satisfy neither equality nor equivalence.

move counter-clockwise, whereas the bottom colors {blue, pink, green, gray} all move clockwise. For Figure 7(b) {red, orange, blue, green} move counter-clockwise, whereas {purple, yellow, pink, gray} move clockwise. Recall also from Table 2 that these are precisely the color divisions that corresponded to Type I and II equivalence patterns.

That patterns 4 and 5 should yield these equivalence patterns is then hardly surprising. By themselves neither pattern has any handedness, so adding a handedness by moving the centers in these particular ways will indeed divide the colors up into exactly the two sets as in Table 2. What is perhaps surprising is that patterns 1 and 3, which also lack handedness, don't behave in exactly the same way, but instead yield equality patterns. So why don't the colors in patterns 1 and 3 care which way their centers moved? We recall that patterns 1 and 3 have an additional symmetry that patterns 4 and 5 lack, namely that they have no preferred axis. As a consequence, the distribution of any single color has an additional three-fold rotational symmetry. This in turn means it doesn't matter which neighbor the center moves to; the resulting combination of corners and center can still always be rotated to match any other color.

That patterns 6-10 yield neither equality nor equivalence is also straightforward to understand. All of these patterns are handed, with the mirror images occurring only between the a and b versions, but all eight colors in any given pattern having the same handedness. If you take any handed pattern though and add something 'left-handed' to four of the colors, and 'right-handed' to the other four, the result will just be a jumble that satisfies neither equality nor equivalence. Finally then, why do patterns 2a,b escape this fate, since they are also handed? Just like patterns 1 and 3, 2a,b also have the additional symmetry of having no preferred axis, and that additional three-fold rotational symmetry is exactly what saves them here, and allows them to continue satisfying equality even when the movement of the centers was deliberately chosen to be in different directions for different colors. We see therefore that all of these symmetry distinctions noted in the article – handed versus non-handed, preferred axis versus no preferred axis – have consequences at this stage as well, when we add in the centers and move them in various ways.

	11a,b	12a,b	13a,d	13b,c	14a,b	15a,b	16a,b	17a,b
Fig. 7(a)	Yes	Yes	–	–	–	–	–	Yes
Fig. 7(b)	–	–	Yes	–	Yes	Yes	–	Yes

Table 4: The results of combining the center arrangements from Figure 7(a,b) with the seven equivalence patterns 11-17. Yes means it satisfies equivalence, '–' means it doesn't.

Table 4 shows the results of adding the Figure 7 centers to the equivalence patterns 11-17. The first point to note is that once the corners have established an equivalence pattern of a given type, adding the centers can never change that type, or change the overall pattern to equality. Adding centers must be consistent with the pre-existing equivalence type, or else the combination of corners and centers simply doesn't satisfy anything. With that in mind, most of the results in Table 4 are again straightforward to understand. In particular, the results for patterns 11, 12, 14, 15, 16 all make perfect sense: (i) 11 and 12, which we recall were Type I equivalence patterns, are consistent with Figure 7(a), (ii) 14 and 15, which are Type II equivalence patterns, are consistent with Figure 7(b), (iii) 16, which is Type III, is not consistent with either of Figures 7(a,b).

The two anomalous patterns are 13 and 17. Even though pattern 13 is Type I, adding the Figure 7(a) centers doesn't work. Perhaps even more surprisingly, adding the Figure 7(b) centers does not work for 13b,c, but it does work for 13a,d! We recall though that 13 was the unusual pattern that had the two preferred axes rather than just one or none. As a result, even though it is indeed a Type I equivalence pattern, the way the colors fit together is slightly different than for the other patterns. For example, the way you have to do the rotations to see the equality within the given color sets is different for 13 than for the other patterns. As a result, center arrangements such as Figure 7(a), which you would think should work, don't necessarily work, whereas center arrangements such as Figure 7(b), which you would think should not work, can sometimes work after all. By the way, can you also construct something similar to Figure 7(b) that works for 13b,c but not for 13a,d? Something like that certainly should exist, since there is nothing fundamentally different between the four versions of pattern 13.

The last item then is pattern 17, which is a Type III equivalence pattern, but somehow still fits together with both of the Figure 7(a,b) center arrangements. The reason in this case is the same as for pattern 2 above. We recall that 17 has those same 'color swirls' as 2 did – and for the same reason, because they both have no preferred axis. It is that additional three-fold rotational symmetry that allows 17 to fit together even with center arrangements that one might otherwise think shouldn't work. So as before for the equality patterns, for the equivalence patterns as well all the previous symmetry considerations play a role here.

As our final example of things one can (or can't) achieve by adding the centers, let's return to Table 3, where we saw that it is possible to start with the equality patterns 4 and 5, and add the centers to produce Type I and II equivalence patterns. Is it possible to also produce Type III equivalence patterns? You should not find it too difficult to arrange the centers in a way which will indeed accomplish this. More intriguingly, suppose we go back to our original six divisions from Figure 10, where only (a,d,e) were relabelled as Types I, II, III, but (b,c,f) didn't yield any equivalence patterns with the colors divided in those ways. Might it be possible after all to construct equivalence patterns of those types, now that we have the centers to play with as well? The answer turns out to be no, which can be worked out without too much effort.

First, we note that it is *only* patterns 4 and 5 that could possibly be turned into any new types of equivalence patterns. As noted above, all of the existing equivalence patterns are immediately excluded, since they can never change their types. The handed equality patterns 2 and 6-10 are also excluded, for the same reason that adding Figure 7 to patterns 6-10 above failed. Once part of the pattern has a handedness, and with all eight colors having the same handedness, then adding anything to it can at most result in another equality pattern (as patterns 2a,b above achieved), or else satisfy neither equality nor equivalence.

Patterns 1 and 3 can be excluded as well, even though they also lack handedness, just like patterns 4 and 5. The problem with both of these is that they are so symmetric that it's impossible to add any handedness to either of them at all. Just try it! Pick any color, for example red, and imagine successively putting the red center into all eight possible positions. Thanks to the three-fold rotational symmetry that both patterns 1 and 3 have, the combined arrangement of the three red patches coming from the corners, together with the red center, simply does not have any handedness, no matter where you put the center. Well, if you can't create a handed pattern for just a single color, you obviously can't create equivalence patterns of any type, since we agreed here that equivalence patterns require the single-color distribution to be handed.

So, how do we show that adding the centers to the equality patterns 4 and 5 can only yield our familiar equivalence types I, II, III, but can't divide the colors in any of these (b,c,f) ways? Let's start with pattern 4, and again take red as our 'reference' color. You then have eight possible choices for where to put that red center. You can easily convince yourself that four of them will yield a combined 'corners + center' pattern that still lacks handedness, and therefore can't yield any equivalence patterns. The other four possible center choices constitute two mirror image pairs. Since equivalence patterns necessarily occur in mirror image pairs anyway, we only have to consider one member of each pair.

Page 31 of the templates document contains two further copies of pattern 4 with the red centers inserted into the relevant positions. Now, what choices do we have for the other seven colors? In particular, which ones do we want to be equal to the red distribution, or mirror images of the red distribution?

Let's start with orange: You should easily be able to convince yourself that orange *must* have the same distribution as red. If you tried to give it the mirror image distribution, its center would occupy the same position that red has already taken. This holds true for either of these templates with the two choices for red, and indeed also for pattern 5, which has exactly the same two choices for the red center.

Similarly, purple and yellow must have the same distribution, so either equal to the red distribution or its mirror image, but both must make the same choice. And yet again for blue and green, and pink and gray. That is, within the four pairs red/orange, purple/yellow, blue/green, pink/gray the distributions must be equal, or else those two colors would clash with regard to their center positions.

So, at this point there are three remaining choices, namely which one of purple/yellow, blue/green or pink/gray goes together with red/orange, and which two have the mirror image distribution. Well, then it's just a matter of comparing with Table 2 and noting that taking purple/yellow to go with red/orange corresponds to Type I, taking blue/green corresponds to Type II, and taking pink/gray corresponds to Type III.

The (b,c,f) divisions simply cannot work, and would inevitably create a clash for some of these four color pairs. It is very nice then that the corners and centers agree so perfectly that equivalence patterns of Types I, II, III are allowed, but these less symmetric (b,c,f) divisions are excluded.

Computational Methodology

As promised, we end with an outline of how to systematically scan through all possible arrangements of the six corner pieces, and pick out the ones that satisfy either equality or equivalence. The focus here is necessarily rather different than in the previous sections, and can get somewhat technical in places, certainly requiring careful attention to detail (as all computing does if you want to get it right). Nevertheless, any readers who are familiar with the idea that an equation of the type $ax+by+cz=d$ represents a plane in three-dimensional (x,y,z) space should be able to follow along.

Specifically, let's consider the plane $x+y+z=6$, shown in Figure 18. The intersection points with the (x,y,z) axes are obviously $(6,0,0)$, $(0,6,0)$ and $(0,0,6)$, as indicated. The midpoints of those line segments are $(0,3,3)$, $(3,0,3)$ and $(3,3,0)$. The coordinates of the three points indicated in red are then $(4,1,1)$, $(1,4,1)$ and $(1,1,4)$, in each case obtained by taking the average of the three surrounding black points. These three red dots are therefore the centers of those little triangles. By the way, the reason for taking the plane as $\ldots = 6$ is to make these numbers come out as integers, and thereby allow the entire calculation to be done in integer arithmetic.

We can clearly do exactly the same thing for all eight planes of the form $\pm x \pm y \pm z = 6$, where the three $\pm$s are taken independently. Next, on the octahedron in its original pattern 1 state, define the x-axis to be along the red/yellow/gray/green corner, the y-axis to be along the red/green/pink/purple corner, and the z-axis to be along the red/purple/orange/yellow corner. It should then be straightforward to verify the validity of Table 5, indicating which plane is associated with each of the eight colors, and therefore what the coordinates of that color's three center points are.

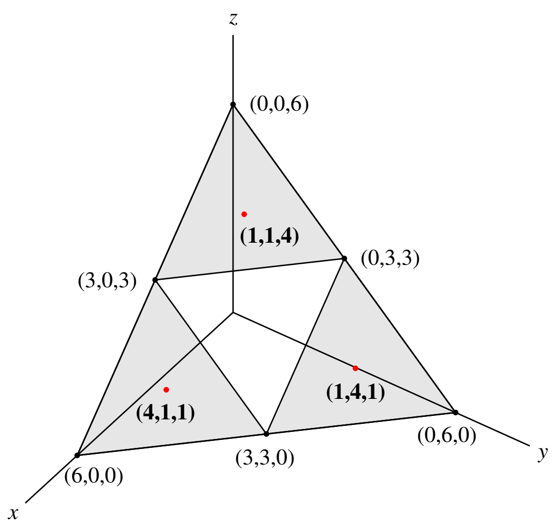

Figure 18: The plane $x+y+z=6$. The gray-shaded triangles will become the corner color patches on the octahedron. As indicated by the red dots, the centers of these triangles can be worked out to be $(4,1,1)$, $(1,4,1)$ and $(1,1,4)$.

Color	Equation of Plane	Three Center Points		
red	$x + y + z = 6$	(4,1,1)	(1,4,1)	(1,1,4)
purple	$-x + y + z = 6$	(-4,1,1)	(-1,4,1)	(-1,1,4)
orange	$-x - y + z = 6$	(-4,-1,1)	(-1,-4,1)	(-1,-1,4)
yellow	$x - y + z = 6$	(4,-1,1)	(1,-4,1)	(1,-1,4)
green	$x + y - z = 6$	(4,1,-1)	(1,4,-1)	(1,1,-4)
pink	$-x + y - z = 6$	(-4,1,-1)	(-1,4,-1)	(-1,1,-4)
blue	$-x - y - z = 6$	(-4,-1,-1)	(-1,-4,-1)	(-1,-1,-4)
gray	$x - y - z = 6$	(4,-1,-1)	(1,-4,-1)	(1,-1,-4)

Table 5: The eight colors on the octahedron in its original configuration, their associated planes, and therefore the center points of their three color patches.

The next step is to recognise that the colors can't be independently rearranged into different positions, but constitute six specific corner pieces that must move as one. We therefore need to rearrange the information in Table 5 to clarify which four colors must go together. This is accomplished in Table 6, which lists six corner pieces, arbitrarily labelled as positions 1-6. For each position, there are the four coordinates that constitute the centers of its four color patches, and the color of that patch, as taken from Table 5. Convince yourself that all the information in Table 6 really is just a rearranged version of the information already in Table 5. Convince yourself also that Table 6 is consistent with your original octahedron, and how the (x, y, z) axes were defined with respect to the eight colors.

The only remaining points to note regarding Table 6 are the order in which the six pieces are listed, and also the order in which the four points for each piece are listed. First, the order in which you list the six corner pieces is completely arbitrary. Next, if you examine the order in which the four colors for each position are listed, and compare with the corresponding template, you will note that they are consistently listed as going counter-clockwise around each corner piece. This must be consistent for all six pieces, but whether you list them clockwise or counter-clockwise, or which color is listed first, is again arbitrary.

Position	Point A	Point B	Point C	Point D
1	(1,1,4)	(-1,1,4)	(-1,-1,4)	(1,-1,4)
2	(1,1,-4)	(1,-1,-4)	(-1,-1,-4)	(-1,1,-4)
3	(1,4,1)	(1,4,-1)	(-1,4,-1)	(-1,4,1)
4	(1,-4,1)	(-1,-4,1)	(-1,-4,-1)	(1,-4,-1)
5	(4,1,1)	(4,-1,1)	(4,-1,-1)	(4,1,-1)
6	(-4,1,1)	(-4,1,-1)	(-4,-1,-1)	(-4,-1,1)

Table 6: The same center points as in Table 5, but now grouped together so that each of the six pieces is on one line of the table. See also page 32 of the templates document.

Position	Point A	Point B	Point C	Point D
1	(1,1,4) ● red	(-1,1,4) ● purple	(-1,-1,4) ● orange	(1,-1,4) ● yellow
2	(1,1,-4) ● gray	(1,-1,-4) ● blue	(-1,-1,-4) ● pink	(-1,1,-4) ● green
3	(1,4,1) ● purple	(1,4,-1) ● red	(-1,4,-1) ● green	(-1,4,1) ● pink
4	(1,-4,1) ● blue	(-1,-4,1) ● gray	(-1,-4,-1) ● yellow	(1,-4,-1) ● orange
5	(4,1,1) ● purple	(4,-1,1) ● pink	(4,-1,-1) ● blue	(4,1,-1) ● orange
6	(-4,1,1) ● red	(-4,1,-1) ● yellow	(-4,-1,-1) ● gray	(-4,-1,1) ● green

Table 7: The six pieces from Table 6 rearranged as indicated.

Now, what can we do with Table 6? Suppose we modify it to produce Table 7. Note first that all 24 coordinates have remained where they are, and it is some of the colors that have moved. That is, the coordinates of all the relevant points really are fixed in our (x, y, z) space, so it makes sense to keep them fixed in the table, and then move the colors to new positions. The six corner pieces have then moved according to:

- The 'top' piece, what we're calling position 1, has remained unchanged.
- The position 2 piece has remained in place, but has been rotated by $-90°$.
- The position 3 piece has remained in place, but has been rotated by $90°$.
- The position 4 piece has again remained in place, but has been rotated by $180°$.
- The position 5 and 6 pieces have switched places, and with their orientations as indicated.

The general rules for moving colors are straightforward: Colors on any given row must always remain together – since these constitute the original six pieces – but can be interchanged with the colors on any other row. The colors on any row must also remain in the same cyclic order – again because the pieces have the colors in specific orders – yielding four possibilities there.

The last page of the templates document has a template for this Table 7 pattern. If you compare each piece with the original octahedron, you should be able to convince yourself that the pieces have indeed moved as indicated above. From your newly constructed octahedron you can also see easily enough that this particular pattern doesn't satisfy either equality or equivalence, but is just some random jumble. How do we get a computer code to recognise that though?

As a first step, let's rearrange the information in Table 7 to produce the first two columns of Table 8. That is, we're essentially doing the reverse of what we previously did in going from Table 5 to Table 6; we're just separately listing three points for each of the eight colors. For the next step, the formal definition of equality would require us to do lots of coordinate rotations, to test whether all eight sets of three points can be mapped into each other. This would obviously be very tedious, and the test for equivalence would be even worse, especially since we don't know in advance which of the six divisions in Figure 10 will actually yield equivalence patterns at all.

Fortunately, there is a shortcut we can use at this point, by using the fact that distances in three-dimensional space are invariant under rotations. That is, if you have two or more points in space, and you rotate them all in exactly the same way, all the distances between them will remain

Color	Three Center Points			Squared Distances		
red	(1,1,4)	(1,4,-1)	(-4,1,1)	34	34	38
purple	(-1,1,4)	(1,4,1)	(4,1,1)	18	22	34
orange	(-1,-1,4)	(1,-4,-1)	(4,1,-1)	34	38	54
yellow	(1,-1,4)	(-1,-4,-1)	(-4,1,-1)	34	38	54
green	(-1,1,-4)	(-1,4,-1)	(-4,-1,1)	18	38	38
pink	(-1,-1,-4)	(-1,4,1)	(4,-1,1)	50	50	50
blue	(1,-1,-4)	(1,-4,1)	(4,-1,-1)	18	22	34
gray	(1,1,-4)	(-1,-4,1)	(-4,-1,-1)	22	38	54

Table 8: The same information as in Table 7, but with the coordinates now grouped by color to make it clearer how the 'equal squared-distances' test is applied. This particular pattern fails, since the eight colors do not have the same three squared distances.

unchanged. So, for each set of three points in Table 8, just compute the three distances between them – or rather, the squared distances, to avoid any need for square-roots, and keep everything as purely integer arithmetic – and then arrange those three squared-distances in order from smallest to largest. The results are shown in the third column in Table 8. We immediately see that the eight colors do not all have the same squared-distances, which is enough to demonstrate that their sets of points couldn't possibly be rotated into one another.

And even better, since distances in three-dimensional space are also invariant under taking mirror images, this 'equal squared-distances' test will automatically also select for equivalence patterns as well, and with any arbitrary division of the colors into the two mirror-image sets. We just have to examine the results by hand afterwards to sort them into equality versus different types of equivalence patterns.

Since this worked so well for this pattern in Tables 7 and 8, and the corresponding template verifies that it is indeed just a random jumble, let's try another example. Table 9 shows a different permutation of the six pieces, and Table 10 again rearranges that information to apply our 'equal squared-distances' test, which this time indicates that this pattern does work. Well, go ahead and tape together the very last of the templates, and verify first of all that the resulting octahedron is indeed consistent with the information in Table 9. Next, compare it with patterns 15a,b. You should be able to recognise that it is the same arrangement, just with a different choice of preferred axis, namely red/yellow/gray/green and blue/orange/purple/pink as the top and bottom pieces.

Finally then, how many permutations of the pieces do we have to check, and how do we consistently keep track of them all? If we return to Table 6, we can always leave the 'position 1' piece unchanged. After all, any pattern can always be rotated to have that particular piece on top, and in that orientation. For the other five pieces there are $5! = 120$ different ways to arrange on which rows they go, and for each choice of rows there are $4^5 = 1024$ orientations. The total number of equivalents of Tables 7 and 9 is thus 122880, which is still a very small number for a computer.

The computer then produced 84 patterns that passed this 'equal squared-distances' test, which

Position	Point A	Point B	Point C	Point D
1	(1,1,4) ● [red]	(-1,1,4) ● [purple]	(-1,-1,4) ● [orange]	(1,-1,4) ● [yellow]
2	(1,1,-4) ● [green]	(1,-1,-4) ● [gray]	(-1,-1,-4) ● [blue]	(-1,1,-4) ● [pink]
3	(1,4,1) ● [gray]	(1,4,-1) ● [yellow]	(-1,4,-1) ● [orange]	(-1,4,1) ● [blue]
4	(1,-4,1) ● [green]	(-1,-4,1) ● [pink]	(-1,-4,-1) ● [purple]	(1,-4,-1) ● [red]
5	(4,1,1) ● [red]	(4,-1,1) ● [yellow]	(4,-1,-1) ● [gray]	(4,1,-1) ● [green]
6	(-4,1,1) ● [blue]	(-4,1,-1) ● [orange]	(-4,-1,-1) ● [purple]	(-4,-1,1) ● [pink]

Table 9: Another rearrangement of the six pieces from Table 6.

Color	Three Center Points			Squared Distances		
red	(1,1,4)	(1,-4,-1)	(4,1,1)	18	38	50
purple	(-1,1,4)	(-1,-4,-1)	(-4,-1,-1)	18	38	50
orange	(-1,-1,4)	(-1,4,-1)	(-4,1,-1)	18	38	50
yellow	(1,-1,4)	(1,4,-1)	(4,-1,1)	18	38	50
green	(1,1,-4)	(1,-4,1)	(4,1,-1)	18	38	50
pink	(-1,1,-4)	(-1,-4,1)	(-4,-1,1)	18	38	50
blue	(-1,-1,-4)	(-1,4,1)	(-4,1,1)	18	38	50
gray	(1,-1,-4)	(1,4,1)	(4,-1,-1)	18	38	50

Table 10: The same information as in Table 9, once again grouped by color to make it easier to see how the 'equal squared-distances' test is applied. This pattern passes the test, since all eight colors have the same three squared distances.

Pattern Nr.	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
Multiplicity	1	2	1	3	3	6	6	6	6	6	6	6	12	6	6	6	2

Table 11: A list of how the 84 raw output patterns are grouped together into the 17 patterns. As indicated in the text, the multiplicities perfectly reflect the symmetries of the patterns.

were further sorted by hand to produce the results in Table 11. At this point, we also encounter our previous symmetries – handed versus non-handed, preferred axis versus no preferred axis – one final time. For example, according to Table 11 pattern 16 has a multiplicity of six, which are exactly the six versions we previously saw in Figure 15. Any pattern with a preferred axis will necessarily have a factor of three in its multiplicity here, or in the case of pattern 13 even a factor of six. Similarly, any handed pattern will necessarily have a factor of two in its multiplicity. All of the multiplicity results are then completely consistent with the symmetries associated with the given patterns – or equivalently, I basically already knew what the symmetries had to be as soon as I saw the multiplicities here, even before studying the finer details of the various patterns.

Well, this really is the end of what turned out to be a great many things to say about patterns on an octahedron! I hope you found it interesting, and learned something about patterns, symmetries in three-dimensional space, and even computational algorithms.