SELF-SIMILAR SELF-SIMILARITY

JOEL DAVID HAMKINS

Abstract. Let me tell a mathematician's tale about symmetry. We begin with playful curiosity about a concrete elementary case—the symmetries of the letters of the alphabet, for instance. Seeking the essence of symmetry, however, we are pushed toward abstraction, to other shapes and higher dimensions. Beyond the geometric figures, we consider the symmetries of an arbitrary mathematical structure—why not the symmetries of the symmetries? And then, of course, we shall have the symmetries of the symmetries of the symmetries, and so on, iterating transfinitely. Amazingly, this process culminates in a sublime self-similar group of symmetries that is its own symmetry group, a self-similar self-similarity.

In the light of symmetry consider a capital letter A.

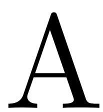

Well, this particular A in this particular font, unfortunately, is not quite perfectly symmetric. The uprights at left and right differ in thickness, for example, and the serif on the right foot is ever so slightly larger than the serif on the left. Let us try to draw a somewhat more symmetric letter A, even though it may be less graceful.

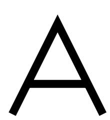

That's better. This A exhibits a vertical line symmetry—the vertical line down the center.

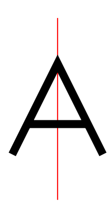

If we reflect the letter across that line—like Alice through the looking glass—it lands perfectly upon itself. We might fold the paper on that line to realize the symmetry.

The letter B also exhibits symmetry.

B

Well, again, this particular B in this particular font is not *perfectly* symmetric—the upper curved half is slightly smaller, and the two curves are not exactly the same shape. But we can try to draw a more symmetric B, if less graceful, so as to exhibit a horizontal line symmetry.

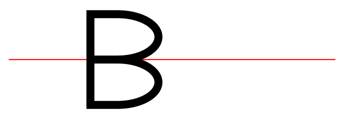

We might similarly consider all the letters of the alphabet, drawing each of them as symmetrically as possible. Get some paper and try!

A B C D E

What symmetries did you find? Some letters, as we have seen, have vertical or horizontal line symmetries. The letter H has both vertical and horizontal line symmetries. The letter S has no line symmetries, but it has a *rotational* symmetry.

S

If you rotate this letter S by 180 degrees about its center, half-way around, it will land precisely upon itself. In elegant fonts, the letter S is often not quite symmetric in this way—the upper part is often gracefully smaller than the lower part. Some letters, such as F, G, and R, seem to have no nontrivial symmetries at all.

Which is the most symmetric letter? The letter X, when drawn with perpendicular lines, exhibits not only vertical and horizontal line symmetries, but also diagonal line symmetries, as well as four-fold rotational symmetry.

X

One particular letter, I claim, can be drawn so as to exhibit *infinitely* many symmetries! Consider the letter O, drawn as a perfect circle.

O

This letter O has line symmetries with respect not just to vertical and horizontal lines, but with respect to any line through the center at all, infinitely many. And it also has infinitely many rotational symmetries, by any angle you like.

To illustrate how the symmetries of a figure form a mathematical system, let us consider the symmetries of a square.

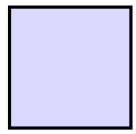

The square exhibits a four-fold rotational symmetry. If we rotate by a quarter turn, the square lands upon itself. We can rotate twice, or three times, but rotating four times brings back the original orientation. This is the same result as with no rotation, the *identity* symmetry, a trivial symmetry leaving the object unchanged. Perhaps the trivial symmetry is hardly a symmetry at all. Yet mathematicians systematically find it useful to regard trivial or degenerate instances of their conceptions as fully valid—every square is also a rectangle; every equilateral triangle is also isosceles; and zero is a number. So let us regard the identity symmetry as a symmetry—it is one of the ways of associating the square rigidly with itself.

The square also exhibits symmetries by reflection, with four lines of symmetry.

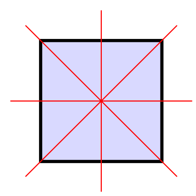

Thus, the square has eight symmetries in all: four rotational symmetries (including the identity symmetry as 0° rotation) and four reflection symmetries.

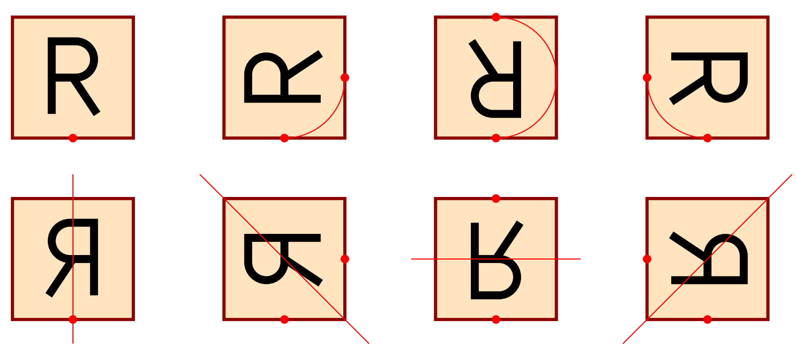

The letter R indicates for each symmetry how it acts upon the square.

Symmetries can be *composed* by performing them one after the other; this is a kind of multiplication operation for symmetries, with the result being another symmetry. Thus, symmetries become dynamic—each is a distinct transformation. If you first reflect on the vertical line and then on the horizontal line, for example, which symmetry have you got?

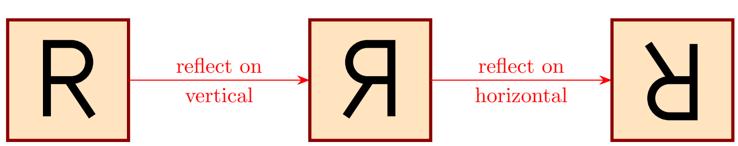

Altogether, this is rotation by 180°, half-way around. The composition of two reflections in the plane is always a rotation, because the mirror orientation, being twice reversed, is ultimately preserved.

Every symmetry admits an *inverse* symmetry, the symmetry that undoes it; composing them is the identity symmetry. The inverse of clockwise rotation is counterclockwise rotation through the same angle; reflections are self-inverse, since performing them twice gives the same result as doing nothing.

When composing symmetries, does the order matter? Yes, indeed it does. If you compose a left quarter turn with a vertical-line reflection in both ways, the end result is not the same.

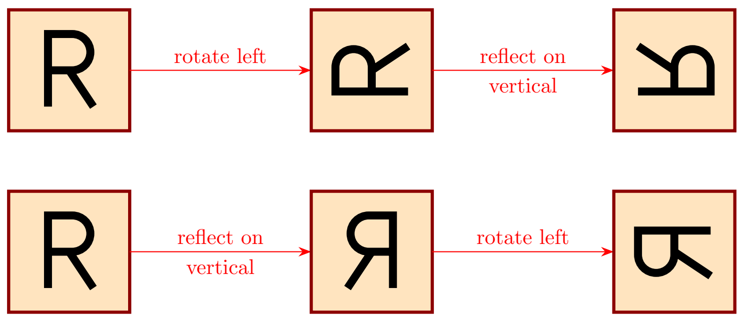

Abstracting to a higher dimension, consider the symmetries of a cube.

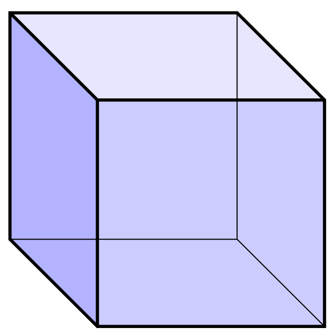

There are numerous rotational symmetries—grab the cube by any face and apply a quarter turn; or grab the cube by a vertex and spin by one-third; or grab the cube by an edge and twist exactly half-way around. There are also numerous mirror-plane reflection symmetries—how many can you find? And there is the *central symmetry*, turning the cube inside out by exchanging each vertex with its opposite through the center. This is neither a rotation nor a planar reflection, it turns out, although it can be realized by composing a rotation with a planar reflection.

How many symmetries of the cube are there? You might be surprised to learn that there are 48 distinct symmetries of the cube. To count them, consider how a symmetry might act upon a particular face. This face must be carried to one of the six faces, and as we observed above, there are eight ways to associate the two squares. Having attached the face, the rest of the symmetry is determined (perhaps requiring reflection), and so there are $6 \times 8 = 48$ many symmetries of the cube altogether.

What is a symmetry? With geometric plane figures, we had rotational and reflective symmetries, and in three dimensions we had the central symmetry. In higher dimensions there are still other kinds of rigid transformations. All of these symmetries are *isometric*, which means that they preserve distances—they do not stretch or compress the figure.

But some figures are insightfully described by nonisometric symmetries. Consider the Sierpinski gasket fractal here, for example. Notice how it appears within itself in scaled-down form—the whole triangular figure appears three times at half size (half the edge size), once in each corner. And it appears nine times at one-quarter size, twenty-seven times at one-eighth size, and so on. There are infinitely many scaled-down copies of the whole fractal inside itself. This is a kind of symmetry, to be sure, a self-similarity, but not by rotation or reflection; because of the scaling, it does not preserve distances. Yet the figure is insightfully described by this self-similarity.

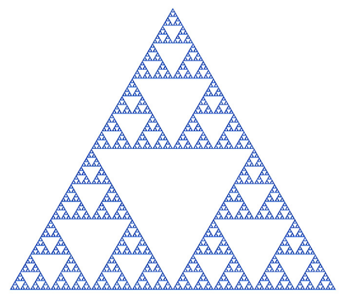

So we are pushed toward a more generous, abstract conception of symmetry. We might consider any structure-preserving transformation of a mathematical object, an isomorphism of the object with itself, as a symmetry of that object. We may consider the symmetries of any mathematical structure at all.

Consider the system of complex numbers $\mathbb{C}$, for example. Complex numbers have the form $a + bi$, where a and b are real numbers and i is the *imaginary unit*, the square root of negative one,

$$i = \sqrt{-1}.$$

We can add complex numbers and multiply them, and altogether the complex numbers form a mathematical structure known as a *field*.

We said that i is the square root of negative one. But suppose I ask, Which one? There are two such complex numbers, since $-i$ also is a square root of negative one, as you can check:

$$(-i)^2 = (-1)^2 i^2 = i^2 = -1.$$

So what had seemed to be the defining property of i is a property that also holds of $-i$. How can we tell them apart?

Indeed, we cannot tell them apart in the complex field. There is nothing you can say about i in the complex number field $\mathbb{C}$ that isn't also true of $-i$. Your i might be my $-i$, for all we know, if we treat the complex numbers strictly as a field. The reason is that there is a symmetry of the complex numbers, an isomorphism of the complex field with itself, that swaps the numbers i and $-i$. This symmetry is called *complex conjugation*, and it associates every complex number with its complex conjugate:

$$a + bi \qquad \longmapsto \qquad a - bi.$$

The numbers i and $-i$ therefore play exactly the same structural role in the complex number field. They are perfectly symmetric copies of each other.

There are many other automorphisms of the complex field—an enormous uncountable infinity of them, $2^{2^{\aleph_0}}$ many—although one uses the axiom of choice to prove this. Every irrational complex number can be moved. And yet, all these symmetries of $\mathbb{C}$, including complex conjugation, are broken by augmenting the complex field with its coordinate structure, using the real and imaginary parts. Thus, the complex plane (as opposed to the mere field) is *rigid*—it has no nontrivial symmetries.

Let us continue our flight toward abstraction. Start with any mathematical structure at all—a geometric figure, a number system, whatever you like—and consider its symmetries. Collectively these constitute the *symmetry group* of your structure. Since this symmetry group is a perfectly good mathematical structure of its own, with composition as the group operation, we may consider *its* symmetries, or in other words, the symmetries of the symmetries of the original structure.

But why stop there? This next symmetry group, after all, also stands on *its* own as a mathematical structure, with *its* symmetry group, the symmetries of the symmetries of the symmetries, and then of course there will be the symmetries of the symmetries of the symmetries of the symmetries, and so on. We may iterate the process as long as we like. In this way, we are led to the *automorphism tower*.

$$G_0 \longrightarrow G_1 \longrightarrow G_2 \longrightarrow G_3 \longrightarrow \cdots$$

We began with the symmetry group G_0 of the original structure, and then each next group is the symmetry group or automorphism group of its predecessor.

It is an elementary fact in group theory that every element of a group generates an inner automorphism of that group by the process of conjugation. Because every group element is thus mapped canonically into its symmetry group, the tower of groups can be seen as building toward a certain limit group G_ω, the direct limit of the system of groups.

$$G_0 \longrightarrow G_1 \longrightarrow G_2 \longrightarrow \cdots \longrightarrow G_\omega$$

And not only that. Because G_ω is a perfectly good group on its own, we may consider its automorphism group, and the automorphism group of *that* group, and so on, thereby continuing the automorphism tower beyond infinity.

$$G_0 \longrightarrow G_1 \longrightarrow G_2 \longrightarrow \cdots \longrightarrow G_\omega \longrightarrow G_{\omega+1} \longrightarrow G_{\omega+2} \longrightarrow \cdots \longrightarrow G_\alpha \longrightarrow \cdots$$

Iterating transfinitely, each next group is the automorphism group of its predecessor, and at limit stages we use the direct limit.

Does it ever stop? Will there ever be a group that is isomorphic to its own automorphism group by the inner-automorphism association we described?

Amazingly, the answer is Yes. The process eventually reaches completion. In my article [Ham98], building on key earlier work of [Tho85], I proved the following:

Theorem. *Every group has a terminating transfinite automorphism tower.*

What I proved is that in every automorphism tower, perhaps very far out in the transfinite part of the tower, there will eventually be a *complete* group, a group for which every automorphism is already realized as an inner automorphism by a distinct group element. Such a group thus exhibits a perfectly self-similar self-similarity. Nothing new is added beyond this group by considering symmetries of symmetries, or symmetries of symmetries of symmetries; one already has them all.

Since every automorphism tower is completed in this way, if you iteratively consider the symmetries of a structure, and then the symmetries of the symmetries, the symmetries of the symmetries of the symmetries, and so on, iterating transfinitely in the natural manner, then you will eventually achieve a complete, sublime group of self-similar self-similarity. You will eventually, perhaps transfinitely, complete the process of symmetry.

References

[Ham98] Joel David Hamkins. "Every group has a terminating transfinite automorphism tower". *Proc. Amer. Math. Soc.* 126.11 (1998), pp. 3223–3226. ISSN: 0002-9939. DOI: 10.1090/S0002-9939-98-04797-2. arXiv:math/9808014[math.LO]. http://jdh.hamkins.org/everygroup/.

[Tho85] Simon Thomas. "The automorphism tower problem". *Proceedings of the American Mathematical Society* 95 (1985), pp. 166–168.

(Joel David Hamkins) Professor of Logic, Oxford University & Sir Peter Strawson Fellow, University College, Oxford

Email address: joeldavid.hamkins@philosophy.ox.ac.uk

URL: http://jdh.hamkins.org