Rubik's cube: finite group, pieces and orientations distinguishable by color, describe as moving edges and moving corners. Edges give an S_{12} for permutation, coerners give an S_8 for permutation. Edges give an $\mathbb{Z}_2^{12}$ for orientation, corners give an $\mathbb{Z}_3^8$ for orientation. But these groups aren't entirely independent.
Look at moving a single face. Get a four-cycle on the edges and a four-cycle on the corners. Both odd. Therefore the parity of the edge permutation is equal to the parity of the corner permutation. Thus we get that the actual permutation group is $A_{12,8} = S_{12} \times S_8/\sim$.
For the facets on the edge pieces, we have a pair of four cycles, so the facet parity is even. Hence we can't flip a single edge, and thus we actually only have an $\mathbb{Z}_2^{11}$ for edge orientation. Similarly, we only have an $\mathbb{Z}_3^7$ for corner orientation.
So in total, we have $|A_{8,12}| \times |\mathbb{Z}_2^{11} \times |\mathbb{Z}_3^7|$ as the size of the Rubik's cube group. This gives 43,252,003,274,489,856,000 possible states.
As it is a group, we can look for things like its center. It has a $\mathbb{Z}_2$ center, with the nontrivial element being called the superflip, where every piece is in the correct position, the corners are oriented correctly, and all of the edges are flipped. We can see this is central because moves that only change orientation are commutative, and as for moves that permute things, the superflip treats each cublet of the same type identically.
Furthermore, it is the only nontrivial state with that property, since if one edge is flipped then all must be flipped, and if one corner is turned by 120 degrees, say, clockwise, we would need all eight corners rotated the same way, but that would violate the three-arity. Hence the superflip is the only nontrivial central element.
The superflip is also interesting in that it is the farthest you can get from being solved. To be more precise, in the half-turn metric, where turning a given face a quarter turn in either direction or a half turn is counted as a single move, Michael Reid in 1995 showed that the superflip takes 20 moves to get to from the solved state, and, by Rokicki, Kociemba, Davidson and Dethridge in 2010, no state takes more than 20 moves to get to.
The proof was basically checking all the possible sequences of 20 moves and showing that they hit everything. Lots of computation. Also, the superflip isn't the only state that takes 20 moves to get to, but it is the nicest one, and also the one that was first proven to take no fewer than 20 moves.
Anyway, that's the Rubik's cube. There are generalizations in a few different directions. The most popular one is to increase the number of subdivisions. So we have cubes with four pieces per edge, and five pieces per edge, and so on. Here we have the interesting new adjustment that not all of the pieces are distinguishable anymore, at least for the standard solid coloring. So now we have nontrivial isotropy, and the action of the group on the cube isn't free anymore. So in theory we have a richer situation here. Actually, the original Rubik's cube group also has isotropy, in that it is possible to spin a single center piece by 180 degrees without changing anything else, or to rotate a pair of centers by 90 degrees; since the Rubik's cube is classically solid-colored, rotating the centers is undetectable. However, since the centers don't move, the sequences that just rotate the centers and leave everything else fixed form a normal subgroup, so we just take the quotient by default. Alternatively, in both the case of the regular Rubik's cube and the higher-order cubes, we can just assume that all of the pieces and their orientations are distinguishable, for instance if we have a picture cube where instead of solid coloring we have pictures

with each patch having a distinguished orientation.
We also have the idea of tacking on more dimensions. So we have four-dimensional Rubik's hypercubes, which have corners, edges, and now mobile faces, where the stickers are three-dimensional cubelets. I've seen implementations that go up to ten dimensions, although these are generally plagued by visualization issues even if the basic solving strategies end up being analogous. Again, this is not what I want to talk about.
And then we also have objects that aren't cubes. The puzzle titled the megaminx is a dodecahedron cut into edges and corners, and again the basic move is to rotate a face. Again we can use the same techniques to talk about this new object. So we count the edges and get 30 edges, and we count the corners and get 20 corners, and so we expect the group to be some subgroup of $S_{30} \times S_{20} \times \mathbb{Z}_2^{30} \times \mathbb{Z}_3^{20}$. Now we look at a single face and turn it. We get that now we have two five cycles for the permutation of pieces. Hence the permutation of the edges must be even and the permutation of the corners must be even independently, whereas with the cube we had that only the product had to be even. This independent evenness holds for any of the puzzles where the number of edges per face is odd.
We still get that you can't flip a single edge or spin a single corner. So now we're in $(A_{30} \times A_{20}) \ltimes (\mathbb{Z}_2^{29} \times \mathbb{Z}_3^{19})$. Again the only nontrivial central element is the superflip. And of course we can do the previous generalizations to this guy as well, increasing the number of pieces per edge, so that a dodecahedron with each edge divided into 5 gives us the Gigaminx, and with 7 we get the Teraminx, and 9 would give us a Petaminx. And I suppose you could move into four dimensions and try to manipulate a 120-cell. I don't think I've ever seen this implemented in code.
But let's think of another way to think about twisting puzzles. Look at the 7x7x7 Rubik's cube. It bulges outward slightly for mechanical reasons, but we can think about going all the way and simply projecting the entire thing onto an enveloping sphere. So all the faces distort into disks on the sphere with circular boundaries, and the inner layers into annuli. Moreover, this gives us a nice way to unify the three-dimensional twisting puzzles, as spheres covered by rotatable patches that have circular boundaries, such that proper rotations preserve the network of circular cuts. And so we see that these twisting puzzles are actually two-dimensional in terms of information content. We don't really care about the inside of the cube or dodecahedron at all, except when trying to make a physical instantiation of these things.
Let's define a Rubik's surface to be a surface covered by disks such that for each disk d_i, there is an integer n_i such that rotation of the disk around its center by $2\pi/n_i$ sends $\bigcup_j \partial d_j$ to itself, i.e. it preserves the network of cuts so that rotations can be composed. We might as well impose the restriction that if d_i is contained in d_j, first that d_i and d_j are concentric and that $n_i = n_j$. We can worry about colorings and such later. Here we just assume that all the pieces are distinguishable. This allows us to talk about cases like the skewb, where we have a cube sliced not parallel to the faces, but across the long diagonals, so that rotations occur around a corner with angle 120 degrees. It is isomorphic to a Rubik's octahedron except the coloring is weird because it thinks it's a cube. We just assume that all the pieces are distinguishable from each other and that all orientations are detectable for corners, edges and faces. If we want we can later impose isotropy to get rid of distinctions.

We've seen now instances of Rubik's surfaces whose underlying manifold is a sphere. What about other surfaces? What about surfaces with positive genus? What about nonorientable surfaces?

Let's look at the nonorientable case first. We have a physical model of a nonorientable Rubik's surface, specifically a Rubik's real projective plane. Normally, to get the real projective plane you take a sphere and identify opposite points. So we take a Rubik's cube and identify opposite cubes. Thus turning the front face clockwise is automatically accompanied by turning the back face counterclockwise, and so on. The result just looks like what you'd get if you made the corners all act as one unit. Alternatively, you can think of it as just moving the middle slices instead of the faces. You can also build such a thing out of the megaminx, or out of an octohedral or icosahedral puzzle. Note: I have never seen a decent, face-turning icosahedral twisting puzzle. Most twisting puzzles that are shaped like icosahedrons are really just other puzzles that have been sanded down into icosahedrons.

Now let's look at orientable, positive genus. The simplest case is the torus. The simplest nontrivial case on the torus is either divide the torus into a pair of triangles or a pair of squares to get dihedrons. But it turns out that dihedrons aren't very interesting, because rotating either face moves all of the mobile pieces, and preserves all of the adjacencies and relative orientations. So really it's just spinning the centers around. Note that this is also the case for dihedrons on the sphere.

So next up is the trihedron of three hexagonal faces on a torus. If we were to take a universal cover of this thing, we'd get the plane tiled with hexagons, with identifications. The tiling of the torus has three faces, six corners, and nine edges. Again, turning a single face gives us an odd permutation of the edges and an odd permutation of the corners, so again every turn of a face is even overall. This tells us pretty quickly that the superflip is impossible here, since that would require nine flips. It also turns out that you can't rotate the corners relative to each other.

So we can move on to the more interesting case, wherein a pair of faces share at most one edge. In this case coloring is sufficient to distinguish all of the pieces. The next simplest torus covering I can think of is a triangular grid with four different faces, but it turns out that this just becomes the tetrahedron; in particular, it turns into a sphere with four degree-2 singular points, but the singular points are hidden as vertices so we don't notice them.

After that, we get to the next case, which is a torus with a square grid and five distnct squares. Note that we can't tile the torus by four squares and have each pair touching only once. In our five-square case, we can look at for the usual features. The tiling gives us five faces, ten edges, and five corners; the corners here have four facets. Let's look at local constraints. Turning a face gives us a four-cycle for edge permutation and a four-cycle for corner permutation, so the $S_{10} \times S_5$ restricts to an $A_{10,5}$. We get a pair of four-cycles for the edge facets, so edge orientation is always even, and we get a quartet of four-cycles for corner facets, so corner orientation has four-arity.

One question that arises naturally is how do we flip even one edge, ignoring everything else? Well, on the normal Rubik's cube the quickest way is to pick a corner adjacent to the edge, and turn the faces adjacent to that corner all in the same direction so that the edge circles around the corner. Because the cube has three faces per edge, the orientation of the edge relative to the corner changes three times and hence when it reaches its original position it's flipped. In the case of this five-tiled torus, we get that we can't do that kind of thing with the corners, because the edge would flip an even number of times. However, we can rotate each of the faces 180 degrees so that the edge moves around the whole torus, and when it gets back to its starting position it's now flipped.

Interestingly, if we compare this to a similar tiling with eight facets, arranged as so, then we get that we can't flip any edge. We can see this by just two-coloring the faces, and then we only have to pay attention to which side of the edge is which color. Since every rotation preserves this, we get that if the edge returns to its original position it must be in its original orientation.

So we can then say that a single edge can be flipped if and only if the graph that is dual to the tiling is not 2-colorable. This is equivalent to saying that a single edge can be flipped only if there is a cycle in the dual graph that has odd length; you just rotate each face corresponding to the vertex in the dual graph along the cycle. So this gives us a global condition that gives us a constraint, as opposed to our previous analyses which only looked at constraints coming from local conditions. So local conditions, i.e. the constraints that we get by looking at the rotation of a single face, are insufficient. There might be more complicated constraints coming from other global considerations, but I don't know what they are.

So far I've been dealing solely with disk structures on manifolds. We don't have to do that. If we were to take, say, a megaminx, i.e. a dodecahedron, then we can identify two commuting generators, in other words two disks that don't intersect. And we'd still get a viable Rubik's structure, but there's no manifold that we could put disks onto to realize this structure with a one-to-one correspondence between disks and faces.

But I have a conjecture, that this is the only bad thing that can happen. We say that M covers R if M is also a Rubik's structure such that R can be obtained from M by identifying disks in M.

Conjecture: If for every cover M of R such that M is a manifold, if for each disk d and each lift $\widehat{d}$ of d, the set of disks that intersect $\widehat{d}$ are the same for all lifts, then R is realizable as a manifold. The heuristic is that rotatability gives us a local manifold structure, and thus we can always cover by some M that can be realized on a manifold. Regularizing the idenfication, which is what that hypothesis does, should prevent indentifications of faces of M that would prevent R from being a manifold as well.

Finally, we can ask when even the group paradigm fails. So let's look at the so-called helicopter cube. Instead of rotating around the faces or corners, this one rotates around the edges. So once you start rotating, you have to rotate it 180 degrees to get back to a cube.

That last statement is a lie. We actually have other things we can do, like turning an edge part of the way, and then turning an adjacent edge, and so on. Note that after doing the first partial turn, certain other edges become fixed. You can actually reach positions where only three of the twelve edges can be turned. So we don't have a group for the helicopter cube; we have a groupoid. And the partial turns make the projection onto a sphere not injective with respect to the surface of the puzzle.

We might hope that the partial turns do not add any new cube-shaped states, but alas, the set of states accessible using only 180 degree turns is a proper subset of the cube shaped states accessible by partial turns. For example, the group of 180 degree turns partitions the center pieces into four orbits, with each piece of the same coloring being in a different orbit. Using partial turns, you get one orbit.