1 Equivalence classes of symmetries

This definition is probably familiar, but it's useful for discussing classification under symmetry.

Definition 1. A set R of ordered pairs from $S \times S$ is called a *relation* on S; two elements a and b are said to be *related* under R, written aRb, if $(a, b) \in R$.

A relation R is said to be *reflexive* if for all $a \in S$, aRa. R is *symmetric* if for all $a, b \in S$, if aRb, then bRa. R is *transitive* if for all a, b, c in S such that aRb and bRc, it follows that aRc. A relation which is reflexive, symmetric, and transitive is known as an *equivalence relation.*

This lends itself to our previous investigation of symmetry as such:

Proposition 1. *For S a set of n-element objects and G a permutation group consisting of permutations of length n, there is a natural relation R_G induced by letting $(a, \pi(a)) \in R_G$ for all $a \in S$ and $\pi \in G$. Then R_G is an equivalence relation on S.*

Proof. Since G is a permutation group, the identity element $e \in G$. By the definition of R_G, it is true that for all $a \in S$, $aR_Ge(a)$, or, since $e(a) = a$, aR_Ga Thus R_G is reflexive.

Now, consider a and b from S such that aR_Gb. By the definition of R_G, this is true because there is a permutation $\pi \in G$ such that $b = \pi(a)$. By the definition of a symmetry group, π^{-1} is also in G, and since it is the inverse of π, it follows that $a = \pi^{-1}b$. Since $\pi^{-1} \in G$ and $b \in S$, $bR_G\pi^{-1}b$, or in other words, bR_Ga. Thus R_G is symmetric.

Lastly, consider a, b, and c in S such that aR_Gb and bR_Gc. By the definition of R_G, there is $\pi \in G$ such that $b = \pi(a)$, and $\sigma \in G$ such that $c = \sigma(b)$. Thus $c = \sigma(\pi(a)) = (\sigma\pi)(a)$. By closure of the permutation group, $\sigma\pi \in G$, and since $a \in S$, it follows that $aR_G(\sigma\pi)(a)$, or in other words, aR_Gc. Thus R_G is transitive. □

This is very useful for our investigation of symmetry because, if G represents some collection of symmetries, then aR_Gb if and only if a and b represent different orientations of the same "object". We thus want to collect elements of S into subsets based on their equivalencies. In fact doing so is an established way of identifying an equivalence relation:

Proposition 2. *The equivalence relations on S are in a one-to-one correspondence with the set-partitions of S.*

Proof. A set-partition of S necessariily induces an equivalence relation: let aRb iff a and b are in the same partition; it is fairly easy to show that this relation definition satisfies the equivalence conditions.

Conversely, if we have an equivalence relation R, we may define a set-partition by letting a and b be in the same partition if and only if aRb. The equivalence conditions guarantee that this is a well-defined procedure: we would only run into inconsistencies if a were required not to be in the same partition as itself (forbidden by reflexivity) or if different established members of a single partition gave different membership-status to an element based on their relations (forbidden by symmetry and transitivity). □

Note, as a corollary, that the number of equivalence relations on a set of size n is the Bell number B_n, since that is the number of set partitions.

Definition 2. The set-partitions induced by an equivalence relation are known as its equivalence classes. The set of equivalence classes of S under the permutation-relation R_G are denoted S/G.

In abstract-algebra parlance, the set of equivalence relations under R_G would be called the *orbits* of S under the group action of G; we shall not use that notation here, but it may be of interest to those using this from an abstract-algebraic standpoint.

This framework allows us to see why some of our symmetry-reduction methods worked before: on several examples we worked previously, G and S were such that, whenever $\pi \neq \sigma$, $\pi a \neq \sigma a$, so that each element a was a member of an equivalence class of size exactly $|G|$, so that $|S/G| = \frac{|S|}{|G|}$.

This sort of equivalence would still be correct, if we "overcount" $|S|$ in a certain manner. Take, for instance, our old example problem of a five-bead necklace in black and white. If we counted only elements of S, each equivalence class would contain different numbers of elements: the necklace BBBBB would be in a class by itself, while BBBBW, BBBWB, BBWBB, BWBBB, and WBBBB would be in an equivalence class together. However, there is an argument to be made that each equivalence class contains 10 *images* of a class representative, e.g., the class $\{BBBBB\}$ might be thoguht to contain the 10 elements $e(BBBBB)$, $r(BBBBB)$, and so forth for every permutation in D_5; likewise, the 5-element class could be thoguht to contain each element twice; for instance, $BBBBW$ might be counted both as $e(BBBBW)$ and as $fr^4(BBBBW)$. If we were to count in this highly unconventional manner, we would find not 32, but 80 elements of S, and then $|S|/|D_5|$ would be 8, as hoped. But how do we codify such an unusual overcounting method? We shall actually do so via expansion of the product $|S/G| \cdot |G|$.

1.1 Burnside's Lemma

S/G is the set of equivalence classes of S under permutations in G; G is a permutation group. We can somewhat trivially represent $|S/G| \cdot |G|$ as the nubmer of pairs of a class A and permutation G:

$$|S/G| \cdot |G| = \sum_{C \in S/G} \sum_{\pi \in G} 1$$

If we arbitrarily choose an element x_C from each equivalence class C, then by unique invertability of permutations, for any pair (C, π) there is exactly one element y of S such that $\pi(y) = x_C$, namely, the value $y = \pi^{-1}(x_C)$. We may thus write the number 1 in an astonishingly ungainly fashion:

$$|S/G| \cdot |G| = \sum_{C \in S/G} \sum_{\pi \in G} |\{y : \pi(y) = x_C\}|$$

which, if we make use of an indicator function, can be made more ungainly too, but ripe for sum-rearrangement:

Definition 3. The indicator function 1_P, where P is a logical proposition, is defined as

$$1_P = \begin{cases} 1 \text{ if } P \text{ is true} \\ 0 \text{ if } P \text{ is false} \end{cases}$$

Then, we may rewrite the above equation as the triple sum:

$$|S/G| \cdot |G| = \sum_{C \in S/G} \sum_{\pi \in G} \sum_{y \in S} 1_{\pi(y) = x_C} = \sum_{y \in S} \sum_{\pi \in G} \sum_{C \in S/G} 1_{\pi(y) = x_C}$$

Noting that $\pi(y)$ will never equal x_C unless y is in the same equivalence class as C, we may see that the innermost sum of our final formulation here has all zero terms except when C is the equivalence class containing y. Armed with this knowledge, we may denote the equivalence class conaining y as $\langle y \rangle$, and remove one of our sums:

$$|S/G| \cdot |G| = \sum_{y \in S} \sum_{\pi \in G} 1_{\pi(y) = x_{\langle y \rangle}}$$

Since y and $x_{\langle y \rangle}$ are in the same equivalence class for all y, there is at least one permutation mapping y to $\langle y \rangle$; we may arbitrarily choose one and call it σ_y, so then we can write the above as:

$$\begin{aligned}
|S/G| \cdot |G| &= \sum_{y \in S} \sum_{\pi \in G} 1_{\pi(y) = \sigma_y(y)} \\
|S/G| \cdot |G| &= \sum_{y \in S} \sum_{\sigma_y^{-1}\pi \in G} 1_{\sigma_y^{-1}\pi(y) = y} \\
|S/G| \cdot |G| &= \sum_{y \in S} \sum_{\pi' \in G} 1_{\pi'(y) = y} \\
|S/G| \cdot |G| &= \sum_{\pi' \in G} \sum_{y \in S} 1_{\pi'(y) = y} = \sum_{\pi \in G} |\{y \in S : \pi(y) = y\}|
\end{aligned}$$

This, together with a new definition, will give us a powerful and efficient new counting method:

Definition 4. The *number of invariants of the permutation π over S*, denoted $\mathrm{Inv}(\pi)$, is equal to the number of $y \in S$ such that $\pi(y) = y$.

Then our calculations above give us:

Lemma 1 (Burnside's Lemma)**.** *For a set S acted upon by the permutation group G,*

$$|S/G| = \frac{1}{|G|} \sum_{\pi \in G} \mathrm{Inv}(G)$$

We may see how this works with several examples, both old and new:

Question 1: *How many 5-bead necklaces can be made with two colors of beads. What about with n colors of beads?*

Answer 1: *Let $S = \{1,2\}^5$, so S consists of 5-tuples of the numbers 1 and 2. The above question asks what $|S/D_5$ is. We shall use Burnside's Lemma to answer this question, considering the 10 elements of D_5 in turn.*

By definition, the identity permutation e satisfies $e(\mathbf{x}) = \mathbf{x}$ for any 5-tuple $\mathbf{x}$. Thus, $\mathrm{Inv}(e) = |S| = 2^5$.

A single rotation maps the first bead to the second, position, the second to the third, and so forth. In order for $\mathbf{x}$ to be invariant under this rotation, it must be the case that $x_1 = x_2$, $x_2 = x_3$, $x_3 = x_4$, $x_4 = x_5$, and $x_5 = x_1$. In other words, all 5 elements of the 5-tuple must be the same. We may select one of the two values and assign it to all five elements; thus, $\mathrm{Inv}(r) = |\{1,2\}| = 2^1$.

The double, triple, and quadruple rotations are similar: for instance, for $\mathbf{x}$ *to be invariant under* r^2*, it is necessary that* $x_1 = x_3 = x_5 = x_2 = x_4$*, again necessitating the same value in all positions of the 5-tuple. Thus* $\mathrm{Inv}(r^2) = \mathrm{Inv}(r^3) = \mathrm{Inv}(r^4) = 2^1/$

Any reflection of the necklace will have the same features, fixing a single element, swapping its neighbors with each other, and swapping the two elements at distance 2 with each other. Thus, in order for $\mathbf{x}$ *to be invariant under, for example, the reflection* $f = (54321)$*, we would have* $x_1 = x_5$ *and* $x_2 = x_4$*, so we have 3 choices: we choose one value for both* x_1 *and* x_5*, one for* x_2 *and* x_4*, and one for* x_3*. There are thus* $|\{1,2\}^3| = 2^3$ *invariants under* f*, and likewise under the other 4 reflections.*

Putting these into Burnside's Lemma, we see that

$$|S/D_5| = \frac{2^5 + 4 \cdot 2^1 + 5 \cdot 2^3}{10} = 8$$

And, if we had n *instead of 2 colors, we would simply have a number other than 2 as the base of all our exponents:*

$$|S/D_5| = \frac{n^5 + 4 \cdot n^1 + 5 \cdot n^3}{10}$$

Note that, as a completely frivolous corollary, we are assured that $n^5 + 5n^3 + 4n$ *is divisible by 10 for all integer* n*, since the above quantity must be an integer.*

We might also, for example, use this same line of argument on 3-dimensional rotations.

Question 2: *How many distinct ways are there to color a the faces of a cube with* n *colors, if cubes are considered to be equivalent if one can be rotated to give the other?*

Answer 2: *A cube has 6 faces, so the underlying set* S *is defined by* $\{1,2,\ldots,n\}^6$*. Now, we must find the invariants of the 24 permutations in the octahedral chiral symmetry group* O*, in order to determine the number of equivalence classes in* S/O*.*

As always, every element of S *is invariant under the identity* e*, so* $\mathrm{Inv}(e) = |S| = n^6$*.*

The three 90° *rotations around the various faces of the cube could be denoted* r_x*,* r_y*, and* r_z*; from an invariant-identification standpoint, these three are identical and we might denote an arbitrary face-axis rotation as* r_f*, also including the inverse rotations* r_x^3*,* r_y^3*, and* r_z^3*. A* 90° *rotation about a pair of faces fixes the two faces being used as axes, but rotates the lateral faces around so that eahc is mapped to its neighbor. Thus, in order for a coloring to be invariant under such a rotation, the 4 lateral faces must be the same color. Thus, an invariant of* r_f *is defined by a choice of three colors: one for each of the two fixed faces, and one for the 4 lateral faces collectively. Thus,* $\mathrm{Inv}(r_f) = \mathrm{Inv}(r_f^3) = n^3$*.*

In contrast, the 180° *rotation about a face-to-face axis will fix the two faces being used as axes, and reorient the 4 lateral faces not so that each is cycled with its neighbors, but so that each is swapped with the opposite face. Thus, an invariant under* r_f^2 *is determined by a choice of four colors: one for each of the fixed faces, and one for each of the two pairs of opposite lateral faces, so that* $\mathrm{Inv}(r_f^2) = n^4$*.*

Moving on to the edge-to-edge rotations, all of them are identical as pertains to invariant identification, so we shall identify them collectively as r_e*.* r_e *leaves no faces fixed and is self-inverse, so it simple swaps 3 pairs of faces. In order to be invariant under* r_e*, a cube must be colored so that each pair of swapped faces in monochromatic; thus, an invariant of* r_e *is determined by a choice of three colors, one for each pair of swapped faces. Thus* $\mathrm{Inv}(r_e) = n^3$*.*

Finally, our vertex-to-vertex rotations, which are again identical as pertains to invariant identification, can be denoted as r_v and r_v^2. These rotate the three faces incident on each end of the axis into each other's positions, so in order to be invariant under either r_v or r_v^3, a cube coloring must color all faces incident on the north pole the same color, and all faces incident on the south pole the same color; thus an invariant is determined by two color choices, and thus $\text{Inv}(r_v) = \text{Inv}(r_v^3) = n^2$.

Assembling all these facts under Burnside's Lemma, noting that r_f, r_e, and r_v are standing in for three, six, and four different axis-choices respectively, we find that

$$|S/O| = \frac{n^6 + 6n^3 + 3n^4 + 6n^3 + 8n^2}{|O|} = \frac{n^6 + 3n^4 + 12n^3 + 8n^2}{24}$$

1.2 The Cycle Index

One notable feature of the above analysis is that what ends up mattering in finding the invariants of a permutation is its number of permutations. To that end, we produce a quantity called the *cycle structure*.

Definition 5. If π can be expressed as a product of disjoint cycles of length $k_1, k_2, k_3, \cdots, k_r$, then π is said to have *cycle structure* $\text{Cyc}(\pi) = x_{k_1} x_{k_2} \cdots x_{k_r}$.

To illustrate cycle structure, we may look at, for instance, the 12 permutations in D_6:

Common name	Mapping representation	Cycle representation	Cycle structure
e	(123456)	(1)(2)(3)(4)(5)(6)	x_1^6
r	(234561)	(123456)	x_6
r^2	(345612)	(135)(246)	x_3^2
r^3	(456123)	(14)(25)(36)	x_2^3
r^4	(561234)	(153)(264)	x_3^2
r^5	(612345)	(165432)	x_6
f	(654321)	(16)(25)(34)	x_2^3
fr	(543216)	(15)(24)(3)(6)	$x_2^2 x_1^2$
fr^2	(432165)	(14)(23)(56)	x_2^3
fr^3	(321654)	(13)(2)(46)(5)	$x_2^2 x_1^2$
fr^4	(216543)	(12)(36)(45)	x_2^3
fr^5	(165432)	(1)(26)(35)(4)	$x_2^2 x_1^2$

We may more succinctly describe D_8 based on cycle structure:

$$\text{Cyc}(D_8) = \sum_{\pi \in D_8} \text{Cyc}(\pi) = x_1^6 + 2x_3^2 + 3x_2^2 x_1^2 + 4x_2^3 + 2x_6$$

which is a comfortable shorthand for saying that D_8 consists of one permutation which has 6 fixed points, 2 permutations which consist of 3 swaps, 3 permutations which consist of 2 swaps and 2 fixed points, 4 permutations which consist of 2 3-cycles, and 2 permutations which are 6-cycles.

There are some easy-to-observe facts about the cycle structure for individual permutations and for permutation groups:

Proposition 3. *If π is a permutation of length n, then* $\mathrm{Cyc}(\pi)\big|_{x_i=x^i} = x^n$.

Proof. If π decomposes into cycles $C_1, C_2, \ldots, C_r$, then each number in $\{1, 2, \ldots, n\}$ appears in exactly one C_i. Thus, $|C_1| + |C_2| + \cdots + |C_r| = n$. By the above decomposition, we know the cycle structure of π to be $x_{|C_1|}x_{|C_2}x_{|C_3|}\cdots x_{|C_r|}$. Thus:

$$\mathrm{Cyc}(\pi)\big|_{x_i=x^i} = x_{|C_1|}x_{|C_2}x_{|C_3|}\cdots x_{|C_r|}\big|_{x_i=x^i} = x^{|C_1|}x^{|C_2}x^{|C_3|}\cdots x^{|C_r|} = x^{|C_1|+|C_2|+\cdots+|C_r|} = x^n$$

□

And as an easy corollary of this result:

Proposition 4. *For any permutation group G,* $\mathrm{Cyc}(G)\big|_{x_i=1} = |G|$.

Proof. Since $1 = 1^n$ for all n, the previous result shows that for any π, $\mathrm{Cyc}(\pi)\big|_{x_i=1} = \mathrm{Cyc}(\pi)\big|_{x_i=1^i} = 1^n = 1$. Thus:

$$\mathrm{Cyc}(G)\big|_{x_i=1} = \sum_{\pi\in G} \mathrm{Cyc}(\pi)\big|_{x_i=1} = \sum_{\pi\in G} 1 = |G|$$

□

An extension of these concepts gives a concise variant of Burnside's lemma:

Proposition 5. *If $S = \{1, 2, \ldots, k\}^n$, then* $\mathrm{Inv}(\pi) = \mathrm{Cyc}(\pi)\big|_{x_i=k}$.

Proof. Let π have cycle decomposition into $C_1C_2C_3\cdots C_r$. We shall show that a vector $\mathbf{x}$ is invariant under π if and only if, for i, j in the same cycle, $x_i = x_j$. By the definition of the cycle form, if i and j are in the same cycle of π, seerated by a distance of ℓ, then π^ℓ is a permutation mapping i to j. However, since $\mathbf{x}$ is invariant under π, $\pi(\mathbf{x}) = \mathbf{x}$, and, applying π multiple times will yield $\pi^\ell(\mathbf{x}) = \mathbf{x}$, and thus, since π^ℓ permutes the ith element to the jth position, it must be the case that $x_i = x_j$.

Since this holds for all i and j within a single cycle of π, it follows that for $C_i = (i_1i_2i_3\cdots i_{|C_i|})$, it must be the case that $x_{i_1} = x_{i_2} = \cdots = x_{i_{|C_i|}}$; we shall define the expression x_{C_i} to be equal to this value. Then, $\mathbf{x}$ is uniquely determined by a selection of the r-tuple $(x_{C_1}, x_{C_2}, \ldots, x_{C_r})$; the invariants of π are thus bijectively mapped to the set $\{1, 2, \ldots, k\}^r$, so $\mathrm{Inv}(\pi) = k^r$.

However, it can also be easily seen that

$$\mathrm{Cyc}(\pi)\big|_{x_i=k} = x_{C_1}x_{C_2}\cdots x_{C_r}\big|_{x_i=k} = kk\cdots k = k^r$$

proving the statement of this proposition. □

Using this idea of the variant, we get an easy formulation of a particular case of Burnside's lemma in the context of cycle structure:

Lemma 2 (Burnside's lemma for unrestricted colorings)**.** *For $S = \{1, 2, 3, \ldots, k\}^n$, if G is a group of length-n permutations, then*

$$|S/G| = \frac{Cyc(G)\big|_{x_i=k}}{|G|}$$