Burnside's Lemma

If G is a finite group of permutations of the elements of a finite set S, then the number of orbits of S under G is given by $\frac{1}{|G|}\sum_{\pi\in G}|inv(\pi)|$.

BACKGROUND:

This result, first proved by Burnside in 1911, is based on the Orbit-Stabilizer Theorem. It addresses the general problem of enumeration of distinct patterns in the presence of symmetry. George Polya generalized the lemma in 1937 with his theory of counting.

KEY DEFINITIONS:

1) A *permutation* is a bijective mapping from a set S to itself.
2) If G is a group of permutations acting on the set S, then for a given $s \in S$, the *stabilizer* of s with respect to G is the subgroup of permutations that leave s fixed. We write this as $stab_G(s)$.
3) If G is a group of permutations acting on the set S, then for a given $s \in S$, the *orbit* of s with respect to G is the subset of elements of S that can be expressed as $\pi(s)$ for some $\pi \in G$. We write this as $orb_G(s)$.
4) If G is a group of permutations acting on the set S, then for a given $\pi \in G$, the *invariant set* of π is the subset of elements of S that are fixed by π, i.e. $\pi(s) = s$ means s is fixed by π. We write this as $inv(\pi)$.

KEY FACTS:

1) The Orbit-Stabilizer Theorem: $|G| = |stab_G(s)| \cdot |orb_G(s)|$

PROOF STRATEGY:

Burnside's Lemma is just a short walk from the Orbit-Stabilizer Theorem. The key is presenting the action of every element of G on every element of S in an array, then flagging the fixed elements, and finally summing the number of hits both row-wise and column-wise. The two sums must be equal, and in one case we get the sum $\sum_{\pi\in G}|inv(\pi)|$, and in the other case we get $\sum_{s\in S}|stab_G(s)|$. The latter sum can be reduced with the aid of the Orbit-Stabilizer Theorem, yielding Burnside's Lemma.

PROOF:

Let us construct an array with the elements of S listed across the top and the elements of G down the side with room for tallies in both directions. It will look something like this:

	s_1	s_2	$\cdots$	s_k	*sum*
π_1	$\circ$	x	x	$\circ$	$\lvert inv(\pi_1)\rvert$
π_2	x	$\circ$	$\circ$	$\circ$	$\lvert inv(\pi_2)\rvert$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
π_n	$\circ$	x	$\circ$	x	$\lvert inv(\pi_n)\rvert$
sum	$\lvert stab_G(s)\rvert$	$\lvert stab_G(s)\rvert$	$\cdots$	$\lvert stab_G(s)\rvert$	

I have placed an x in the array to indicate a "hit" whenever $\pi_j(s_i) = s_i$. The number of hits in a row is the size of the invariant set for the permutation at the beginning of the row, and the number of hits in a column is the number of permutations that fix a particular element at the top of the column, namely the order of its stabilizer. Since the overall number of hits in the array is fixed, the two sums $\sum_{\pi\in G}|inv(\pi)|$ and $\sum_{s\in S}|stab_G(s)|$ must be equal.

Now by the Orbit-Stabilizer Theorem, $\sum_{s\in S}|stab_G(s)| = \sum_{s\in S}\frac{|G|}{|orb_G(s)|}$, so we have $\sum_{s\in S}\frac{|G|}{|orb_G(s)|} = \sum_{\pi\in G}|inv(\pi)|$. This gives us $\sum_{s\in S}\frac{1}{|orb_G(s)|} = \frac{1}{|G|}\sum_{\pi\in G}|inv(\pi)|$.

Suppose $orb_G(s) = \{\sigma_1, \sigma_2, \cdots \sigma_m\}$, so that $|orb_G(s)| = m$. Whenever the index of $\sum_{s\in S}\frac{1}{|orb_G(s)|}$ reaches an element of $orb_G(s)$, it contributes the term $\frac{1}{|orb_G(s)|} = \frac{1}{m}$. Since there are m such elements in the overall index set, the total contribution to the sum by the elements in $orb_G(s)$ will be $m \bullet \frac{1}{m} = 1$. In other words each orbit (recall the orbits partition the set S) contributes a 1 to the sum $\sum_{s\in S}\frac{1}{|orb_G(s)|}$. So the total number of orbits of S under the action of G is evidently$\sum_{s\in S}\frac{1}{|orb_G(s)|} = \frac{1}{|G|}\sum_{\pi\in G}|inv(\pi)|$, as required.

APPLICATION:

A *coloring* χ (Greek $\chi\rho o\mu o\varsigma$, color) is a mapping from the finite set S of objects to be colored to the finite set P, consisting of the "pallette" of colors. This is an amusingly fanciful way of thinking about the rather dry notion of an arbitrary function between two finite sets, much in the spirit of the physicists who have their red, green, and blue quarks. Colorings had their provenance in the effort to establish the Four-Color Theorem, and they pop up in many combinatorial settings. If we have a group G of permutations that shuffle the objects of S, then those permutations would likewise shuffle the possible colorings of S. So a given $\pi \in G$ can be regarded as a mapping from the set of P-colorings of S back to itself. More mathematically, there is a group of induced maps G^* that form an action on the set of colorings P^S. Each $\pi \in G$ gives rise to (induces) a companion mapping $\pi^* \in G^*$ that determines what happens to the available colorings whenever the underlying objects in S are permuted by π. It should be clear that G^* is a group with the same order as G.

Just as the elements of S can be put into equivalence classes on the basis of whether they are in the same orbit under the action of G or not, colorings can be put into equivalence classes depending on whether they are in the same orbit under the action of G^* or not. Orbits of colorings are called *patterns*. Two colorings χ_1 and χ_2 are equivalent, or represent the same pattern, if there is a $\pi \in G$ such that for all $s \in S$ we have $\chi_1(s) = \chi_2(\pi(s))$. For example, if $\chi_1(a) = red$, $\chi_1(b) = white$, and $\chi_1(c) = blue$, then if the permutation $\pi = (a \quad b \quad c) \in G$, χ_1 and χ_2 are equivalent provided

$\chi_2(a) = blue$, $\chi_2(b) = red$, and $\chi_2(c) = white$. You can see that $red = \chi_1(a) = \chi_2(\pi(a)) = \chi_2(b) = red$, and so forth. We write $\pi^*(\chi_1) = \chi_2$ to signify this situation.

If we have a symmetrical object with a coloring, we can apply our theory to enumerate the number of distinct patterns possible for the object. Suppose we have a tic-tac-toe grid and we are interested in finding all the distinct ways the cells can be filled in with X's, O's, and blanks. We are only interested in the number of distinct "boards", or patterns...some of them may represent incomplete or impossible tic-tac-toe games. The 3 by 3 grid has the cyclic group of order 4 as its symmetry group, and not the dihedral group D_4 because we regard mirror image games as distinct (lefty vs righty), however rotating a board through a multiple of 90 degrees does not count as a distinct pattern. Let us think about the invariant sets of the four rotations through 0, 90, 180, and 270 degrees, respectively. The identity permutation, rotation through 0 degrees, has 3^9 fixed points, corresponding to the number of ways nine cells can be colored with three colors, X, O, and blank. Recall the number of functions from a k-set to an n-set is n choices k times, or n^k. The 90 degree rotation takes each corner cell into the next one counterclockwise, so to be invariant, the coloring must have the same color for all corner cells. Similarly, the middle cells in the periphery must all have the same color, however it need not be the common color of the corner cells. Finally, the center cell of the grid can be any color. We thus have $3 \times 3 \times 3 = 27$ colorings left fixed by the 90 degree rotation, and by symmetry, the 270 degree rotation. For the 180 degree rotation, diametrically opposite cells in the grid must be the same color for the coloring to be an invariant. This means 9 possible choices for coloring the corner cells, another 9 for the middle peripheral cells, which need not match the corner cells, and 3 choices again for the center cell, all of which yields 243 possible colorings that are fixed by the 180 degree rotation.

Thus $\frac{1}{|G|}\sum_{\pi \in G}|inv(\pi)| = \frac{19683 + 27 + 243 + 27}{4} = 4995$ is the number of distinct patterns achievable with the tic-tac-toe grid, where two boards are considered to be the same if one can be rotated to yield the other. Of course, the all blank game and the all X game are among these, so if we wanted to make a more realistic estimate of the total number of distinct games possible, we would need to eliminate the impossible games. Polya's Theorem allows estimates of this nature to be refined by adding a level of detail that would quantify such impossible games so that they could be deducted from the overall estimate.