1 Restricted symmetry-enumerations

So far we've built what is essentially a magic bullet to answer the question: "if we freely paint an object with k colors, how many distinct colorings are there subject to symmetry group G?" Recall that our answer to that question was:

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

However, if the painting is restricted in some way, then we'll need a different formula; recall that our central conceit $\mathrm{Inv}(\pi) = \mathrm{Cyc}(\pi)\big|_{x_i=k}$ was based on the idea that each cycle can be freely painted any of k colors, which might not be true if we restrict our color choices. By analogy to our old friend, the twelvefold way, we can look at two particular coloring restrictions.

1.1 Using each color no more than once

If we restrict ourselves to using each color no more than once, we start by reducing the size of the non-symmetry-reduced S: instead of $|S| = k^n$, we now have $|S| = \frac{k!}{(n-k)!}$ (and if the number of features n of the object being colored exceeds the number of colors, no such colorings exist, and $|S| = 0$). Otherwise, however, this case will be easy to analyze, because most permutations have no invariants in S.

Proposition 1. *If S consists of the colorings of an object not using the same color twice, and π is a nontrivial permutation, then* $\mathrm{Inv}(\pi) = 0$.

Proof. Since π is nontrivial, there must be some $i \neq j$ such that $\pi(i) = j$. Suppose $\mathbf{x}$ is an invariant of π. Then it must be the case that $x_i = x_j$, since i is mapped onto j by the permutation π. But since $i \neq j$, it would thus follow that two distinct elements of $\mathbf{x}$ are the same color, which would preclude $\mathbf{x}$ being in S; thus our premise that $\mathbf{x}$ is an invariant of π is contradictory, and there can be no suce $\mathbf{x}$. □

Therefore, Burnside's Lemma is easy to apply in this circumstance, since all but one element of the sum of the invariants will be zero:

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

Note that, in this special case, Burnside's Lemma reduces down to the simple symmetry-reduction techniques we learned very early in this course! Examples of these cases were given as the bridge-table and first beaded-necklace problem from October 22nd's notes.

1.2 Using each color no less than once

Here we have a case which is in a number of ways unfamiliar, and not quite as simplifiable as either the free selection or the unique-color case. We shall start with an example to work our way towards a general theory:

Question 1: *We have 3 colors which we shall use to color the faces of an octahedron. If we must use each color at least once, how many ways (up to rotation-equivalence) are there to do this? What if we had n colors?*

Answer 1: *We start by defining our underlying set S, which are the colorings of the faces of an octahedron using each color at least once, and* not *worrying about multiple orientations of the same coloring. We can basically view the colorings as surjective mappings from the 8 faces to the 3 colors, which we know can be enumerated using the twelvefold way, so* $|S| = 3!S(8,3) = 5796$.

Now, let's consider the invariants of the 24 rotational symmetries of the octagon. We konw from previous investigation that there are only 5 cases we really need to look at: e, r_v (appearing 6 times), r_v^2 (appearing 3 times), r_e (appearing 6 times), and r_f (appearing 8 times).

Of course, $\text{Inv}(e) = |S| = 3!S(8,3)$. *As we saw previously, invariance under r_v requires that the entire northern hemisphere be monochromatic, and that the entire southern hemisphere be monochromatic; in other words, it is dictated by a choice of 2 colors, which must satisfy the original condition that all three of the colors are used. In other words,* $\text{Inv}(r_v) = 3!S(2,3) = 0$. *Now, looking at r_v^2, we know it to divide the faces into 4 swap-pairs, and in order to be invariant under r_v^2, a coloring must use the same color for each face in a pair; thus, invariants of r_v^2 are dictated by a choice of 4 colors, subject to our initial condition that a coloring use all 3 colors, so* $\text{Inv}(r_v^2) = 3!S(4,3) = 36$. *Similarly, r_e matches the faces into four swapped pairs, subject to the same analysis, so* $\text{Inv}(r_e) = 3!S(4,3) = 36$ *as well. Lastly, r_f consists of two fixed points and 3 cycles of length 3; each of these 4 cycles must be monochromatic, so we are again selecting 4 colors using each of the 3 colors possible, so that* $\text{Inv}(r_f) = 3!S(4,3) = 36$.

Using Burnside's lemma, it will thus follow that

$$|S/O| = \frac{3!S(8,3) + 6 \cdot 3!S(2,3) + 3 \cdot 3!S(4,3) + 6 \cdot 3!S(4,3) + 8 \cdot 3!S(4,3)}{24} = 267$$

Using n colors instead of 3, *the only major change would be replacement of most of the* 3*s with n, giving:*

$$|S/O| = \frac{n!S(8,n) + 17 \cdot n!S(4,n) + 6 \cdot n!S(2,n)}{24}$$

We can summarize our process here with the following chart:

Permutation	*Multiplicity*	*Cycle index*	*Number of cycles*	*Number of invariants*
e	*1*	x_1^8	*8*	$n!S(8,n)$
r_v	*6*	x_4^2	*2*	$n!S(2,n)$
r_v^2	*3*	x_2^4	*4*	$n!S(4,n)$
r_e	*6*	x_2^4	*4*	$n!S(4,n)$
r_f	*8*	$x_1^2x_3^2$	*4*	$n!S(4,n)$

The example above gives a pretty good idea of where we stand in developing a general rule:

Proposition 2. *If S consists of all surjective mappings of n objects onto k colors, then* $\text{Inv}(\pi) = k!S(r,k)$, *where r is the number of cycles in π.*

Proof. We saw in a proof last week that, if π has cycle decomposition into $C_1C_2C_3 \cdots C_r$, then a vector $\mathbf{x}$ is invariant under π if and only if, for i,j in the same cycle, $x_i = x_j$, and thus that $\mathbf{x}$ is uniquely determined by a selection of the r-tuple $(x_{C_1}, x_{C_2}, \ldots, x_{C_r})$. However, $\mathbf{x}$ contains all elements of $\{1,2,\ldots,k\}$ if and only if this $r-$tuple contains all elements of $\{1,2,\ldots,k\}$; thus an invariant member of S is uniquely determined by a *surjective* mapping from $\{C_1, C_2, \ldots, C_r\}$ to $\{1,2,\ldots,k\}$, of which there are $k!S(r,k)$ in total. □

This unfortunately doesn't have a simple concise representation as a replacement for the x_i in $\text{Cyc}(\pi)$ the way the free color choice did, but it's still computationally fairly straightforward.

1.3 Specific color-control

Suppose we want to get very specific about the colors we're using: for instance, one might ask how many 9-bead necklaces there are using three green beads, four white beads, and two black beads. This one's pretty tricky! We could do it by restricting S to the set of $\binom{9}{3,4,2}$ elements and then working out the invariants under each permutation, but that's a somewhat ad-hoc solution. What gives us a better result is searching for a *pattern inventory*, which is somewhat akin to a generating function:

Definition 1. The *pattern inventory* for a structure being colored k colors is the polynomial $\sum q_{a_1,a_2,\ldots,a_k} c_1^{a_1} c_2^{a_2} c_3^{a_3} \cdots c_k^{a_k}$, where $q_{a_1,a_2,\ldots,a_k}$ is the number of ways to color the structure with a_i elements of color i.

So if we could get the pattern inventory for the 9-bead necklace with 3 colors, we could answer our question by teasing out the coefficient of $c_1^3 c_2^3 c_3^2$.

What we essentially want to do is modify Burnside's Lemma so that instead of merely counting invariants, we multiply them by appropriate polynomials to produce a pattern inventory.

Question 2: *What is the pattern inventory for coloring a 9-bead necklace with green, black, and white beads?*

Answer 2: *Let's consider the invariants of e identified with regard to the beads they use. Any element of S is invariant under e, so we have free color selection for nine beads. For each bead, we can choose green, black, or white, so the polynomial corresponding to a single bead is $(g^1 + b^1 + w^1)$; since we perform this action 9 tiems, keeping a running count, we'll get $(g+b+w)^9$ as a polynomial representing the pattern inventory of invariants under e.*

Now, let us consider invariants under the one-bead rotation r (or r^2, r^4, r^5, r^7, or r^8; they all have the same cycle structure). This permutation consists of an enormous 9-bead rotation, so all 9 beads must be the same color. Thus, the patterns invariant under one-bead rotation are quite simple: either all 9 beads are green, all 9 are white, or all 9 are black, represented by the pattern inventory $g^9 + b^9 + w^9$.

Now consider r^3 (or r^6); this has three 3-bead cycles, so we have 3 choices of colors, each contributing 3 beads of its chosen color. A single selection thus has inventory $g^3 + b^3 + w^3$; the inventory describing the results of all three selections is $(g^3 + b^3 + w^3)^3$.

Similarly, the flip f or any flip fr^i has one fixed point, and 4 swaps, giving pattern inventory $(g + b + w)(g^2 + b^2 + w^2)^4$.

Our pattern inventory for the 9-bead necklace is thus, using Burnside's lemma:

$$\frac{(g + b + w)^9 + 6(g^9 + b^9 + w^9) + 2(g^3 + b^3 + w^3)^3 + 9(g + b + w)(g^2 + b^2 + w^2)^4}{18}$$

Question 3: *How many 9-bead necklaces are there using three green beads, four white beads, and two black beads?*

Answer 3: *We could use a computer algebra system, or hand computation on the above polynomial to find the coefficient of $g^3w^4b^2$. Looking term-by-term, $(g + b + w)^9$ has a term $\binom{9}{3,4,2} g^3w^4b^2$; $2(g^3 + b^3 + w^3)^3$ has no $g^3w^4b^2$ term (since the exponents are not multiples of 3), and $9(g + b + w)(g^2 + b^2 + w^2)^4$ has the term $9\binom{4}{1,2,1} g^3w^4b^2$ (the first factor must contribute a g, in order for* *the remaining terms to be of parity which could appear in* $(g^2 + b^2 + w^2)^4$*). Thus, we have the coefficient:*

$$\frac{\binom{9}{3,4,2} + 9\binom{4}{1,2,1}}{18} = 76$$

We can see a generalization arising here. We can start with an incredibly traditional, essentially generating-functional observiation:

Proposition 3. *If* S *consists of the colorings of an* n*-feature object with* k *colors, the pattern inventory of* S *is* $(c_1 + c_2 + c_3 + \cdots + c_k)^n$*.*

Proof. This is a straightforward extension of the theory built behind generating functions: the coloring of a single feature can proceed in any of k different ways, with each of the ways incrementing the count of a particular color; thus, coloring a single feature is associated with the pattern inventory $c_1^1 + c_2^1 + \cdots + c_k$; coloring n features has pattern inventory equal to the products of the individual pattern inventories of the features, or $(c_1 + c_2 + c_3 + \cdots + c_k)^n$. □

But now we can generalize that idea to a generating-functional variant of Burnside's Lemma:

Theorem 1 (The Pólya Enumeration)**.** *If* S *consists of all the* k*-colorings of an object with* n *features, and* G *is a permutation group acting on* S*, then the pattern inventory for* S/G *is*

$$\frac{\mathrm{Cyc}(G)\big|_{x_i = c_1^i + c_2^i + \cdots + c_k^i}}{|G|}$$

Proof. For brevity, given a vector $\mathbf{a} = (a_1, a_2, \ldots, a_k)$ of color-quantities (so that $a_1 + a_2 + \cdots + a_k = n$), let $c^{\mathbf{a}}$ denote the monomial $c_1^{a_1} c_2^{a_2} \cdots c_k^{a_k}$. Then a pattern inventory on S/G is simply the sum $\sum_{\mathbf{a}} q_{\mathbf{a}} c^{\mathbf{a}}$, where $q_{\mathbf{a}}$ is the number of elements of S/G which have colors occurring in the numbers dictated by $\mathbf{a}$. Let $S_{\mathbf{a}}$ be the subset of S consisting of those colorings which use colors in quantities determined by $\mathbf{a}$. Note that $S_{\mathbf{a}}$ is closed under permutation, since permuting the features of an object has no effect on the number of colors used. Thus, we can reasonably talk about $S_{\mathbf{a}}/G$, ad in particular, from the definition of $q_{\mathbf{a}}$ above, it is clear that $q_{\mathbf{a}} = |S_{\mathbf{a}}/G$, so the pattern inventory we seek will be simply:

$$\sum_{\mathbf{a}} |S_{\mathbf{a}}/G| c^{\mathbf{a}}$$

to which we can apply Burnside's Lemma:

$$\sum_{\mathbf{a}} |S_{\mathbf{a}}/G| c^{\mathbf{a}} = \sum_{\mathbf{a}} \frac{1}{|G|} \sum_{\pi \in G} \mathrm{Inv}_{S_{\mathbf{a}}}(\pi) c^{\mathbf{a}} = \frac{1}{|G|} \sum_{\pi \in G} \left(\sum_{\mathbf{a}} \mathrm{Inv}_{S_{\mathbf{a}}}(\pi) c^{\mathbf{a}} \right)$$

Note that the parenthesized expression above is itself a pattern inventory, since it has as the coefficient for $c^{\mathbf{a}}$ the number of configurations using color-quantities $\mathbf{a}$ which are invariant under π, it is exactly the pattern inventory for the set of colorings invariant under π.

We have previously enumerated the colorings which are invariant under π; determining their pattern inventory is a minor tweak on the same procedure. Recall that in that proof, we demonstrated that, for π with cycle decomposition $C_1 C_2 C_3 \cdots C_r$, a coloring is invariant iff it is monochromatic on each cycle C_i. Thus, a coloring which is invariant under π is uniquely determined by a choice of color for each of $C_1, C_2, \ldots, C_r$. We may thus find a pattern inventory for the invariants under π by multiplying the pattern inventories associated with coloring each cycle. In coloring a cycle C_i, we choose a color for the entire cycle, and color all $|C_i|$ features that color; this can be done in any of k ways, but each choice of a color increases the count of that color by the number of features colored, namely $|C_i|$. Thus, the pattern inventory associated with selecting a color for the cycle C_i is $c_1^{|C_i|} + c_2^{|C_i|} + \cdots + c_k^{|C_i|}$. Assembling the pattenr inventories for all the cycles, we find that the pattern inventory for invariants under π as a whole is:

$$\begin{aligned}\sum_{\mathbf{a}} \mathrm{Inv}_{S_{\mathbf{a}}}(\pi) c^{\mathbf{a}} &= (c_1^{|C_1|} + c_2^{|C_1|} + \cdots + c_k^{|C_1|})(c_1^{|C_2|} + c_2^{|C_2|} + \cdots + c_k^{|C_2|}) \cdots (c_1^{|C_r|} + c_2^{|C_r|} + \cdots + c_k^{|C_r|}) \\ &= x_{C_1} x_{C_2} \cdots x_{C_r}\Big|_{x_i = c_1^i + c_2^i + \cdots + c_k^i} \\ &= \mathrm{Cyc}(\pi)\Big|_{x_i = c_1^i + c_2^i + \cdots + c_k^i}\end{aligned}$$

and thus, the pattern inventory for S/G is

$$\frac{1}{|G|} \sum_{\pi \in G} \mathrm{Cyc}(\pi)\Big|_{x_i = c_1^i + c_2^i + \cdots + c_k^i} = \frac{\mathrm{Cyc}(G)\Big|_{x_i = c_1^i + c_2^i + \cdots + c_k^i}}{|G|}$$

□

This means everything we did earlier with tedious exhaustion on the necklace can be done easily if we happen to know $\mathrm{Cyc}(G)$.

Question 4: *What is the pattern inventory for coloring the faces of the tetrahedron in four colors?*

Answer 4: *Let us start by classifying the permutations in the tetrahedral rotation group and identifying their cycle index to get* $\mathrm{Cyc}(T)$*. There are 12 permutations: one is the identity, eight are rotations of opposite face-vertex pairs, and three are rotations about edge-pairs. The face-vertex rotation can be easily seen to give one fixed point and a cycle of length three, while the edge-pair rotations swap pairs of faces adgacent to the opposite edges:*

Permutation	*Multiplicity*	*Cycle index*
e	*1*	x_1^4
r_{vf}	*8*	$x_1 x_3$
r_e	*3*	x_2^2

So $\mathrm{Cyc}(T) = x_1^4 + 8x_1x_3 + 3x_2^2$*. Thus our pattern inventory is:*

$$(c_1 + c_2 + c_3 + c_4)^4 + 8(c_1 + c_2 + c_3 + c_4)(c_1^3 + c_2^3 + c_3^3 + c_4^3) + 3(c_1^2 + c_2^2 + c_3^2 + c_4^2)^2$$

which could be expanded algebraically to find the number of instances of individual patterns.

Note that the pattern inventory data *includes* the total count. If we were to evaluate with all the $c_i = 1$, the coefficients alone would drop out, giving the total count — and, in this case, the Pólya enumeration would simplify to the specialized Burnside's Lemma form $\frac{\mathrm{Cyc}(G)\big|_{x_i=k}}{|G|}$

2 Structured Combinatorial Objects

We are now ready to move into a wider field of combinatorial study. Up until now, pretty much everything we have looked at falls into the general category of "enumerative combinatorics": that is, systems in which the only interesting element is the count. Even the structures we've been looking at are only really of enumerative interest: we've looked mostly at sets, with occasional digressions into ordered n-tuples and multisets.

We're interested now in *structures*: sets together with relationships among the elements of those sets. We will spend enxt semester on *graphs*, which are one of the richer combinatorial objects, but for now we will look at a simpler sort of object.

2.1 Posets

We have encountered relations already, as well as the properties of reflexivity and transitivity. We have also encountered symmetry, which is a useful property of relations which describe equivalence, but for relations which describe some manner of unequal comparison, we more often have a different property:

Definition 2. A relation R on a set S is *antisymmetric* if, for all $a, b \in S$, if aRb and bRa, then $a = b$. Alternatively, we can describe this by saying that for $a \neq b \in S$, at most one of aRb and bRa is true.

With this property, we have a class of properties which a reasonable relation used for comparison should have:

Definition 3. A relation $\preceq$ is a *partial ordering* if it is reflexive, antisymmetric, and transitive. An ordered pair $(S, \preceq)$ consisting of a set and an ordered relation on that set is called a *partially ordered set*, or *poset*.

We normally look at finite posets when studying combinatorics, such as $(\{1, 2, 3, 4, 5\}, \leq)$, or $(\{1, 2, 4, 5, 6, 8, 9, 12\}, |)$ (note that $a \mid b$ denotes "a divides b"), or $(\mathbf{P}(\{a, b, c\}), \subseteq)$, where $\mathbf{P}(S)$ is the set of all subsets of S. Infinite posets certainly exist, such as $(\mathbb{R}, \leq)$ or $(\mathbb{Z}^+, |)$, but we will generally not consider them too much.

After introducing the concept of posets, we now must simply produce several definitions describing their properties. We first note that antisymmetry allows tree possibilities for unequal elements: two of these, $a \prec b$ and $b \prec a$ have common-sense meanigns to us, but what of the case where neither of those is true?

Definition 4. If $a \not\prec b$ and $b \not\prec a$, then a and b are called *incomparable*. If $a \prec b$ or $b \prec a$, then a and b are called *comparable*

Definition 5. A poset in which all elements are comparable is called a *totally ordered set* or a *chain*. A poset in which any two distinct elements are incomparable is called an *antichain*.

One problem which we will look at later is the question of finding large subsets of a poset which are chains or antichains. For instance, considering $(\{1, 2, 4, 5, 6, 8, 9, 12\}, |)$, the subset $\{1, 2, 4, 8\}$ is a chain, while $\{5, 6, 8, 9\}$ is an antichain.

Transitivity also gives us an angle on posets: namely, since the ordering is noncyclic, there are prospects of largest or smallest elements of a set. However, we have two distinct concepts of "largest" thanks to incomparability.

Definition 6. a is *maximal* (*minimal*) in $(S, \prec)$ if there is no $b \neq a$ in S such that $a \prec b$ ($b \prec a$).

Definition 7. a is the *greatest element* (*least element*) of $(S, \prec)$ if for all $b \in S$, $b \prec a$ ($a \prec b$).