These notes are a continuation of "Bijective Proofs and Catalan Numbers" from last time.

Thus far, everything that we have said about sets works in the infinite setting. Now, we will discuss one useful result that only makes sense in the finite setting:

Proposition. Suppose we have a function $f : A \to B$, where A and B are finite sets. Then:

(a) If f is 1-to-1, then $\#(A) \leq \#(B)$;

(b) If f is onto, then $\#(A) \geq \#(B)$;

(c) If $\#(A) = \#(B)$, then f is 1-to-1 $\iff$ f is onto.

A 1-to-1 correspondence is just a function that is both 1-to-1 and onto. Thus (c) states that if $\#(A) = \#(B)$ and we want to prove that f is a 1-to-1 correspondence, it suffices to prove that f is *either* 1-to-1 or onto.

Proof. For finite sets A, B, recall that there is a 1-to-1 correspondence $A \longleftrightarrow B$ if and only if $\#(A) = \#(B)$.

(a) Since f is 1-to-1, it restricts to a 1-to-1 correspondence $A \longleftrightarrow f(A)$ (see previous notes). But $f(A) \subset B$ and the size of a subset is clearly $\leq$ to the size of the whole set, so $\#(A) = \#\big(f(A)\big) \leq \#(B)$.

Suppose, in addition, that $\#(A) = \#(B)$. Then $\#\big(f(A)\big) = \#(A) = \#(B)$, so $f(A)$ is a subset of B with the same number of elements as B. This can only happen if $f(A) = B$, so f is onto.

(b) Since f is onto, for any $b \in B$, we can find some $a \in A$ with $f(a) = b$. Choosing such an a for every b, we get a new function $g : B \to A$ such that $f\big(g(b)\big) = b$ for every $b \in B$. (This may **not** be an inverse, as we may not have $g\big(f(a)\big) = a$ for all $a \in A$). But g is 1-to-1, because $g(b_1) = g(b_2)$ implies that

$$b_1 = f\big(g(b_1)\big) = f\big(g(b_2)\big) = b_2.$$

By part (a), we now see that $\#(B) \leq \#(A)$, as desired.

Suppose, in addition, that $\#(A) = \#(B)$. Then we proved in part (a) that a 1-to-1 function $g : B \to A$ must be onto. Suppose that $f(a_1) = f(a_2)$. Since since g is onto, there exist $b_1, b_2 \in B$ with $g(b_1) = a_1$ and $g(b_2) = a_2$. Since $f\big(g(b)\big) = b$ for any $b \in B$, we now have

$$b_1 = f\big(g(b_1)\big) = f(a_1) = f(a_2) = f\big(g(b_2)\big) = b_2.$$

Therefore $a_1 = g(b_1) = g(b_2) = a_2$, which proves that f is 1-to-1.

We have proved part (c) in two parts, following part (a) and part (b). □

We now deal with n-permutations, which are just the various ways to re-order the set $[n] = \{1, \ldots, n\}$. In other words, we can view a permutation as a 1-to-1 correspondence $\sigma : [n] \to [n]$. Permutations are more concrete when written out as a sequence of numbers, but the function σ simply tells you what number $\sigma(i)$ is in the i^{th} position of your sequence. For example, we can write the 4-permutation 1432 as a function:

$$\begin{array}{cccc} 1 & 2 & 3 & 4 \\ \downarrow & \downarrow & \downarrow & \downarrow \\ 1 & 4 & 3 & 2 \end{array}$$

Another great way to represent a permutation σ is pictorially. We take an $n \times n$ grid and in the i^{th} row, we fill in the square in the $\sigma(i)^{\text{th}}$ column. As an example, see the figure below, which represents $\sigma = 1432$. Such a grid represents a permutation if and only if there is exactly one filled square in each row and column.

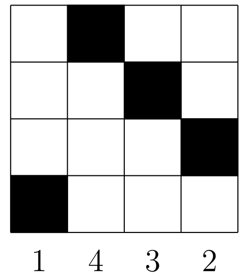

The depiction of permutations in this way allows us to investigate geometric properties of permutations. (For example, we can look for permutations that are rotationally symmetric.) The key notion in permutation patterns is that of containment. If σ is an n-permutation and τ is a k-permutation with $k \leq n$, we say that σ **contains** τ if the diagram of τ can be formed from the diagram of σ by deleting the row and column corresponding to a filled square, or by iterating this operation. For example, note that 32451 contains 231:

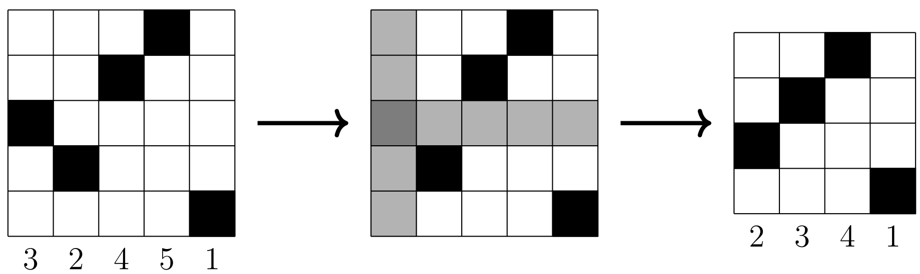

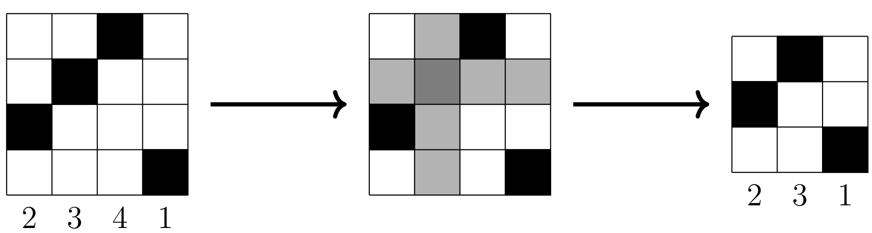

Morally, containment treats a permutation as a sequence and deletes some elements. But we must adjust the resulting sequence to consist of the numbers $\{1, \ldots, k\}$, while preserving their relative order. The above example can then be viewed as deleting 3 and 4, to get $32451 \longrightarrow 251$. Since a 3-permutation cannot contain the number 5, we instead take 231. But the relative ordering of the entries of this list remains unchanged.

The opposite of containment is avoidance: we say that σ **avoids** τ if σ does not contain τ. Then for any permutation τ, we can try to count the number a_k of τ-avoiding k-permutations, for any k. This is a fruitful way to compare permutations, because two permutations that look different may yield the same sequences.

Definition. Two permutations σ and τ are said to **Wilf-equivalent** if for every integer $k \geq 1$, the numbers of σ-avoiding and τ-avoiding k-permutations are equal. A fundamental and difficult question in the study of permutation patterns is the determination of Wilf-equivalence classes of permutations.

Exercise. Fix an n-permutation σ. How many k-permutations avoid σ when $k < n$? What about $k = n$? Using this, prove that Wilf-equivalent permutations must have the same "length."

Now, the question is to classify Wilf-equivalence classes among n-permutations. For $n = 3$, this has long been known. For $n = 4$, the results were discovered by Zvezdelina Stankova. I don't know as much about the higher cases, but I believe that Stankova and West resolved the cases of $5 \le n \le 7$, while all cases $n \ge 8$ remain open. In these notes, we *begin* the process of showing that all permutations in S_3 are Wilf-equivalent, by showing that the number of 231-avoid n-permutations is the n^{th} Catalan number C_n. The steps are:

1. Let S_n denote the set of n-permutations and let L_n denote the set of integer sequences $(a_1, \dots, a_n)$ with $0 \le a_k < k$ for each $k = 1, \dots, n$. Let $A \subset S_n$ denote the set of 231-avoiding permutations.

2. We define $F : S_n \to L_n$ and show that F is onto. Since $\#(S_n) = n! = \#(L_n)$, we see that F is 1-to-1.

3. We will show that $F(A_n) \subset L_n$ is precisely the set of non-decreasing functions in L_n. We saw last time that there are C_n such sequences, so $\#(A) = \#\big(F(A)\big) = C_n$, as desired.

We now commence with the actual proof.

Proof. Step 1 only has definitions, so we begin by defining the function $F : S_n \to L_n$ mentioned in Step 2. For any n-permutation σ and any $k = 1, \dots, n$, we define $F_k(\sigma)$ to be the number of integers i such that $0 \le i < k$ and $\sigma(i) < \sigma(k)$. This clearly gives $0 \le F_k(\sigma) < k$, so the sequence $\big(F_1(\sigma), \dots, F_n(\sigma)\big)$ is in L_n. Let $F(\sigma)$ denote this sequence. We can view $F_k(\sigma)$ as the number of filled squares that are before-and-down from the k^{th} square, as illustrated below for $F(32451) = (0, 0, 2, 3, 0)$.

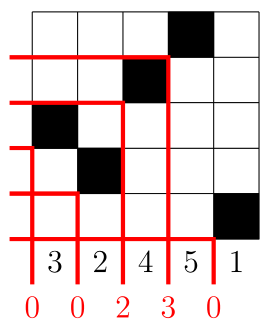

The last element of $F(\sigma)$ has a rather simple relationship to σ. Recall that $F_n(\sigma)$ is the number of values $1 \le i < n$ such that $\sigma(i) < \sigma(n)$. Since σ is a permutation, each value $1, 2, \dots, \sigma(n) - 1$ is attained as $\sigma(i)$ for some $i < n$. Thus $F_n(\sigma) = \sigma(n) - 1$. This will be useful below, as we prove that F is onto.

2. Consider any sequence $(a_1, \dots, a_n) \in L_n$. Then $(a_1, \dots, a_{n-1}) \in L_{n-1}$ and we assume by induction that we have already proved the result for all smaller n. Thus, we may assume that $F : S_{n-1} \to L_{n-1}$ is onto, so there is some $\tau \in S_{n-1}$ with $F(\tau) = (a_1, \dots, a_{n-1})$. We define a function $\sigma : [n] \to [n]$ by

$$\sigma(n) = a_n + 1 \qquad \text{and} \qquad k < n \implies \sigma(k) = \begin{cases} \tau(k), & \tau(k) < a_n + 1 \\ \tau(k) + 1, & \tau(k) \ge a_n + 1 \end{cases}$$

We need $\sigma(n) = a_n + 1$ to get $F_n(\sigma) = \sigma(n) - 1$. Intuitively, we transformed

$$\big(\tau(1), \dots, \tau(n-1)\big) \longrightarrow \big(\sigma(1), \dots, \sigma(n-1)\big)$$

so that the order is not changed, but values $\ge a_n + 1$ are shifted up to leave a gap for $\sigma(n) = a_n + 1$. (In particular, note that σ contains τ.) This gives an explanation for why σ is a permutation and why

$F_k(\sigma) = F_k(\tau) = a_k$ for all $k = 1, \dots, n-1$. Rather than writing out all of the details, we will just illustrate this inductive process for the sequence $(0, 0, 2, 1, 2) \in L_n$. This will take the following form:

$$F(1) = (0) \longrightarrow F(21) = (0,0) \longrightarrow F(213) = (0,0,2)$$
$$\longrightarrow F(3142) = (0,0,2,1) \longrightarrow F(41253) = (0,0,2,1,2).$$

The base case of the induction is $F(1) = (0)$, as these are the *only* elements of S_1 and L_1, respectively. We build up from here by the geometric process described above.

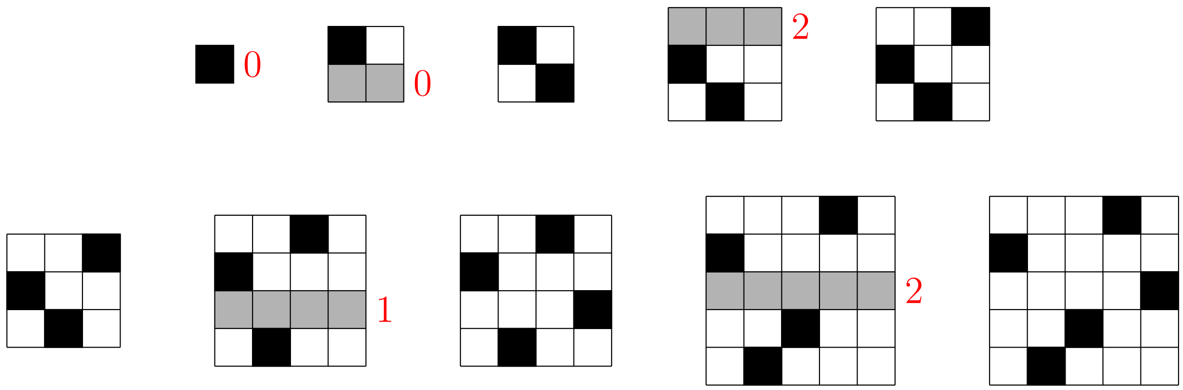

The general procedure follows a similar pattern to this. We have now considered an arbitrary sequence $(a_1, \dots, a_n) \in L_n$ and found some $\sigma \in S_n$ with $F(\sigma) = (a_1, \dots, a_n)$. This proves that F is onto.

3. For any $k = 1, \dots, n$, we can write I_k for the set of all $1 \leq i < k$ with $\sigma(i) < \sigma(k)$. Our definition of F can be written as $F_k(\sigma) = \#(I_k)$. This is useful in showing that σ is 231-avoiding if and only if $F(\sigma)$ is non-decreasing. We prove this in two steps.

   (a) First suppose that $F(\sigma)$ decreases somewhere, i.e. there exist some $j < k$ such that $F_j(\sigma) > F_k(\sigma)$. Then $\#(I_j) > \#(I_k)$, so there exists some integer $i \in I_j$ that is not in I_k. This means that $i < j$ and $\sigma(i) < \sigma(j)$, but $\sigma(j) > \sigma(k)$. Therefore $i < j < k$ and $\sigma(k) < \sigma(i) < \sigma(j)$, which is equivalent to the assertion that σ contains 231 (at the positions i, j and k).

   (b) Conversely, suppose that σ contains 231. Then there are some $i < j < k$ with $\sigma(k) < \sigma(i) < \sigma(j)$. Fix i and j; let k be as small as possible with $k > j$ and $\sigma(j) < \sigma(i)$. If we had $h \in I_k$ and $h > j$, then we would also have $\sigma(h) < \sigma(k) < \sigma(i)$, contradicting the minimality of k. Thus if $h \in I_k$, then $h \leq j$ and $\sigma(h) < \sigma(k) < \sigma(j)$. This proves that $I_k \subset I_j$. Notice that $I_k \neq I_j$, because $i \in I_j$ but $i \notin I_k$. Therefore, we must have $F_k(\sigma) = \#(I_k) < \#(I_j) = F_j(\sigma)$, so F decreases.

   We have now shown that σ contains 231 if and only if $F(\sigma)$ decreases somewhere, which is equivalent σ being 231-avoiding if and only if $F(\sigma)$ is non-decreasing.

This concludes the proof, as all of the other details are contained in the outline above. □

Sadly, this is where we will leave off. The remaining details in the Wilf-equivalence of all 3-permutations can be found in many papers by experts in the field, including this one. Alternatively, there is a blog post that I wrote last year, which narrates my first exposure to this material and gives full proofs of the results. This area of math is not at all my specialty, so please take my blatant self-promotion with a grain of salt. Thank you so much for listening/reading.