The Hidden and Surprising Structure of Ordered Lists

Lara Pudwell

Some mathematical ideas can seem so straightforward at a first glance that we take them for granted and move along to meet more complicated ideas. But sometimes a shift in perspective on a familiar idea opens up a whole new world to explore. For instance, a *permutation* is a list of objects where order matters. There are $n! = n \cdot (n-1) \cdots 2 \cdot 1$ permutations of n objects because there are n choices for the first object, n – 1 choices for the second object, and so on. For example, the ordered lists 123, 132, 213, 231, 312, and 321 are the $3! = 3 \cdot 2 \cdot 1 = 6$ permutations of the numbers 1, 2, and 3. Friendly definition, friendly computation. What more can be said?

Permutations in Permutations

We can make permutations more visual by *graphing* them. To graph a permutation, we read its entries one at a time, and if the entry in position i is the object j, we plot the point (i,j) in the xy-plane. Figure 1 displays the graphs of all six permutations of 1, 2, and 3.

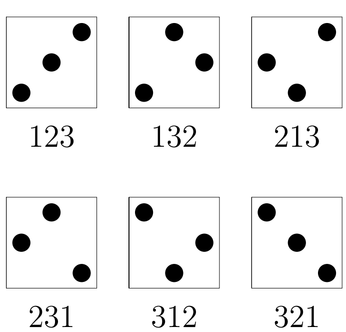

Figure 1. Graphs of the permutations of three objects.

The graphs of permutations provide a different point of view. In list form, one might be tempted to group permutations according to the same first number, but in figure 1, we might notice that the graphs of the permutations 123 and 321 look linear while the graphs of the other four permutations look triangular.

Not only do these graphs give a more visual interpretation than lists, they allow for an interesting twist: finding smaller permutations *inside* of larger ones.

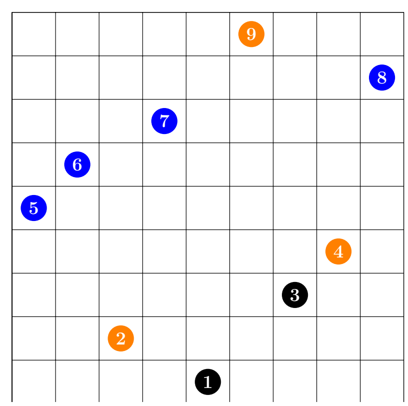

Figure 2. The graph of 562719348 with certain patterns highlighted.

Let's look at a larger example like the graph of 562719348 in figure 2. Suppose we're not interested in the entire permutation, but only focus on some of the dots—such as the 2, the 9, and the 4 (highlighted in orange). Notice that the least of the three numbers, 2, comes first, followed by the largest of the three numbers, 9, followed by the middle of the three numbers, 4, just like in the graph of 132. These three orange dots form a stretched-out version of the graph of 132 inside of the larger permutation, so we say 562719348 *contains* 132 as a *pattern*. Similarly, the dots 5, 6, 7, and 8 (shown in blue) are in increasing order within 562719348; these blue dots form a 1234 pattern within the larger permutation. On the other hand, it's not possible to find four dots in decreasing order in figure 2 (like in the graph of 4321), so 562719348 *avoids* the pattern 4321.

Let's Count!

While it's a great game of mathematical hide and seek to start with a large permutation and look for smaller patterns inside of it, it's actually an interesting research question to start with a small permutation and ask "How many permutations of length n contain my pattern?" or, equivalently, "How many permutations of length n avoid my pattern?"

Once we answer one of these questions, we know the answer to the other one because every permutation of length n either contains pattern p or avoids pattern p. Thanks to this fact, the number of permutations avoiding p plus the number of permutations containing p always adds up to the total number of permutations, or $n!$ for length n. These pattern containment and pattern avoidance definitions first appeared in computer science work in the 1960s, when Donald Knuth showed that permutations sortable in a particular algorithm were exactly the permutations that avoid 231 (*The Art of Computer Programming: Volume 1*, Addison Wesley, 1968). Consequently, mathematicians who work in this area tend to focus on "How many permutations of length n avoid my pattern?"

Let's start with small patterns and see how far we can get! How many permutations of length n avoid the pattern 12? There is only one permutation of length 1, and it's too small to contain a 12 pattern, so there's one 12-avoiding permutation of length 1. There are two permutations of length 2. One of them is 12, but the other, 21, avoids 12, so there's one 12-avoiding permutation of length 2.

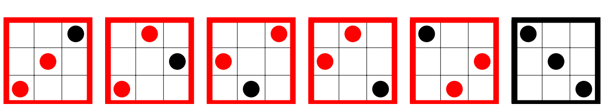

Figure 3. Examples of the pattern 12 contained in all but one of the permutations of length 3.

In figure 3, the six permutations of length 3 are shown with copies of 12 highlighted in red. There is exactly one 12-avoiding permutation in this figure, and it's 321. This is no accident. A copy of 12 is a smaller value followed by a larger value. If a permutation avoids 12, its dots must be in decreasing order. There is exactly one permutation of each length n that avoids 12. Similarly, if a permutation avoids 21, the dots in its graph must be in increasing order, so there is exactly one permutation of each length n that avoids 21 as well.

Scaling up and trying to avoid the pattern 132 is more work, but we get a lovely sequence. Let C_n be the number of 132-avoiding permutations of length n. The first value is $C_0 = 1$ as there is only one way to draw zero dots, and that way avoids 132 (it avoids everything). Similarly, the next two values of the sequence are $C_1 = 1$ and $C_2 = 2$ because any permutation of length 1 or length 2 is not long enough to have three dots and thus also avoids 132. For larger permutations, consider the graph shown in figure 4.

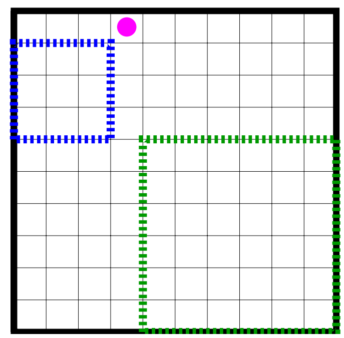

Figure 4. The graph of a 132-avoiding permutation.

The purple dot in figure 4 represents the position of the number n. First, we can observe that "n appears somewhere", so perhaps the purple dot belongs in the first column, in the last column, or somewhere in between. Being a little more precise, let's say n appears in column i with $1 \leq i \leq n$. That means there are $i - 1$ dots to the left of the purple dot and $n - i$ dots to its right. If any of the dots on the left were lower than any of the dots on the right, we could build a 132 pattern by using one dot on the left, one dot on the right, and the purple dot in between them; so every dot to the left of the purple dot must be higher than every dot to the right of the purple dot. In other words, all the other dots in the graph of the 132-avoiding permutation must be inside of the blue box, containing $i - 1$ dots, and the green box, containing $n - i$ dots, shown in figure 4. As long as we fill in the blue box with one of the C_{i-1} 132-avoiding permutations of length $i - 1$ and fill in the green box with one of the C_{n-i} 132-avoiding permutations of length $n - i$, then, voila! We've produced a larger 132-avoiding permutation of length n. Considering all possible options for i—the position for the largest object n—we get the recurrence

$$C_n = C_0C_{n-1} + C_1C_{n-2} + C_2C_{n-3} + \cdots + C_{n-2}C_1 + C_{n-1}C_0.$$

We can check that $C_2 = C_0C_1 + C_1C_0 = 1 + 1 = 2$, as expected. We can also compute $C_3 = 5$, $C_4 = 14$, and $C_5 = 42$.

In general, it turns out that a closed formula for C_n is given by

$$C_n = \frac{1}{n+1}\binom{2n}{n},$$

and this is actually quite a famous sequence known as the *Catalan numbers*. They appear in literally hundreds of counting problems in addition to our exploration of pattern-avoiding permutations. (Richard Stanley, *Catalan Numbers*, Cambridge University Press, 2015). For example, while it is more work, it can be shown that the number of permutations of length n avoiding p where p is your favorite permutation of length 3 is also C_n.

While we're focused on patterns of length 3, let's see if we can avoid multiple patterns at the same time: How many permutations of length n avoid all three of 132, 213, and 123 at the same time? Take a minute and gather some data on your own. How many permutations of length 1 avoid these patterns? How about of length 2? Length 0? Length 3? Can you find a pattern? We'll let F_n be the number of these permutations. As before, $F_0 = 1$, $F_1 = 1$, and $F_2 = 2$ because the permutations in question are too short to contain the length 3 patterns. We now have some extra restrictions compared to the picture in figure 4. In fact, our permutations must now fit into one of the two cases shown in figure 5.

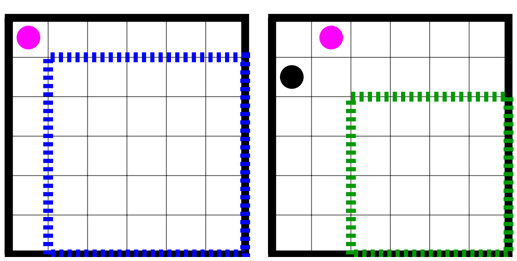

Figure 5. Graphs of permutations avoiding the patterns 132, 213, and 123 simultaneously.

Why? The first observation is that the digit n must appear in column 1 or column 2. If not, then the numbers in the first two columns together with n must form either form a 213 pattern or a 123 pattern. If n is in the second position, we know that avoiding 132 forces the dot before the n to be larger than all the dots after the n. Thus, if n is second, then $n - 1$ must be first.

Thus, if n is in position 1, we can fill in the blue box, shown on the left in figure 5, with one of the F_{n-1} pattern-avoiding permutations of length $n - 1$ to get a pattern-avoiding permutation of length n. If n is in position two (so that $n - 1$ is in position one), we fill in the green box on the right of figure 5 with one of the F_{n-2} pattern-avoiding permutations of length $n - 2$ to get a pattern-avoiding permutation of length n. In other words, $F_n = F_{n-1} + F_{n-2}$. This recurrence, along with the initial conditions, demonstrates that the number of permutations of length n avoiding the patterns 132, 213, and 123 is given by the Fibonacci numbers—yet another famous sequence!

Zooming Out

We have counted the permutations avoiding 12 (1 of each length), avoiding 132 (C_n of length n), and avoiding 132, 213, and 123 simultaneously (F_n of length n). Of course, we could avoid more patterns at the same time, or we could avoid longer patterns. It turns out that although there are 4! = 24 permutation patterns of length 4, if you ask "How many permutations of length n avoid p?" when p is a pattern of length 4, only three distinct counting sequences arise. The number of 1234-avoiders is 1, 1, 2, 6, 23, 103, 513, 2761, … the number of 1342-avoiders is 1, 1, 2, 6, 23, 103, 512, 2740, …, and the number of 1324-avoiders is 1, 1, 2, 6, 23, 103, 513, 2762, … (Miklos Bóna, *Combinatorics of Permutations*, Chapman & Hall, 2004).

The first two of these sequences were explored by researchers in the 1990s, and there are known formulas for expressing the nth term in each of these sequences; however, to this day no one know the precise number of 1324-avoiding permutations of length 1,000,000 or more. We have computed upper and lower bounds on how fast this sequence grows, but its exact values for large values of n remain a mystery. A simple definition, given in pictures, leads to interesting and challenging counting questions quite quickly.

Even more broadly, this isn't just an enumeration game. As mentioned earlier, counting permutations that avoid patterns originally grew out of questions asked by computer scientists. Beyond counting, pattern-avoiding permutations have shown up in biology, physics, and chemistry questions and more.

***Lara Pudwell** is a professor of mathematics and statistics at Valparaiso University. She is also the executive director of MathPath, a national summer camp for middle schoolers who love mathematics.*

A talk including the content of this paper is at
www.youtube.com/watch?v=B7DPo9YQTgw.