Templates for Pattern Avoidance

Yonah BIERS-ARIEL

Abstract:

1. One approach to finding the sizes of permutation avoidance classes is to construct easily enumerated sets and then see if these sets avoid any interesting patterns. In this section, we develop a method of generating sets of permutations using **templates** which both avoid certain patterns, and grow quickly as the lengths of the permutations increase. We will define two kinds of templates, but first will try to motivate their definition with a proof of the well-known fact that the number of permutations of length n which avoid the pattern 132, a quantity which we will call B_n, is equal to C_n, then n^{th} Catalan number.

Theorem 1 The number of 132-avoiding permutations of length n is given by C_n.

Proof: The proof is by induction. When $n = 0$, it is clear that $B_n = 1$, so suppose that $B_m = C_m$ for all $m < n$. Consider a length-n permutation π, and suppose that n appears in position i. If π avoids 132, it follows that the $i-1$ numbers which proceed n must all be greater than all the $n-i$ numbers which follow n, and, moreover, the prefix of π formed by the first $i-1$ numbers and the suffix formed by the last $n-i$ numbers must both avoid 132 themselves. Conversely, if these two conditions are met, then π avoids 132. Any instance of 132 cannot have the 1 and the 2 on opposite sides of the number n because every number preceding n is greater than every number following it, but any instance of 132 also cannot have the 1 and the 2 on the same side of n because both the prefix preceding n and the suffix following n avoid 132 (and, obviously, neither 1 nor 2 can be represented by n). It follows by induction that $B_n = \sum_{i=1}^{n} B_{i-1} \cdot B_{n-i}$ for all $n \geq 0$; since B_n has the same initial condition as C_n and follows the same recurrence, we conclude that $B_n = C_n$ for all $n \geq 0$.

In this proof, we showed that every 132-avoiding permutation of length n has the form LnS where L and S are 132-avoiding permutations such that every number in L is larger than every number in S. We will generalize this idea in the following definition.

Definition 2: A **template** of length $n \geq 1$ is a pair of strings P and B of length n. We require that P be a permutation of length n and B be a binary string of length n. We will denote the i^{th} element of P by p_i and the i^{th} element of B by b_i.

For every positive integer n and template $T = (P, B)$, we define a set of permutations of length n, which we will call $R_{n,T}$, as follows. First, $R_{0,T}$ is the empty string and $R_{1,T} = \{1\}$ regardless of T. Then, $R_{n,T}$ is the set of permutations π of length n which can be divided into subwords (i.e. strings of consecutive elements of π) called $W_1, ..., W_t$ (with $t = |P| = |B|$) such that if $p_i > p_j$, then every of W_i greater than every element of W_j. Moreover, we require that each W_i of length l be an element of $U_{l,T}$, and, if $B_i = 0$, then W_i must have exactly one element. If these conditions are met, we say that $W_1, ..., W_t$ **fit** the template T, so a permutation of length n is an element of

$R_{n,T}$ if it can be decomposed into subwords which fit T. We now provide an example of the set of a template.

Example 3: Let $T = (231, 101)$; then $R_{1,T} = \{1\}, R_{2,T} = \{12, 21\}$, and $R_{3,T} = \{123, 213, 231, 312, 321\}$. To find the elements of $R_{3,T}$ we consider a permutation π of length 3 and divide it up into subwords $W1, W2, W3$. We know that $W2$ is the string 3, and so we can choose $W1 \in R_{2,T}$ and $W3$ empty, $W3 \in R_{2,T}$ and $W1$ empty, or $W1, W3 \in R_{1,T}$. Because $|R_{2,T}| = 2$, each of the first two options gives two distinct permutations in $R_{3,T}$ (123, 213, 312, and 321), while the last option gives one permutation (231). Note that $R_{3,T}$ is exactly the set of length 3 permutations which avoid 132. In fact $R_{n,T}$ is the set of length n permutations which avoid 132; this fact can be checked by reviewing the proof of Theorem 1. Therefore, considering sets corresponding to templates does generalize the argument of Theorem 1.

Once we begin looking at permutations with length greater than 3 it becomes much harder (and likely impossible) to find templates which produce entire pattern avoidance classes. However, it is not too difficult to find templates which produce only permutations avoiding some set of patterns, which is to say subsets of pattern avoidance classes. Therefore, looking at templates lets us find lower bounds on the size of certain avoidance classes. The following proposition shows an application of this method.

Proposition 4: Let Q_n be the set of all permutations of n which avoid every element of $\{2143, 2413, 3142\}$ and let $q_n = |Q_n|$. Then, if the sequence $(r_n)_{n=0}^{\infty}$ is defined by $r_0 = r_1 = 1$, and $r_n = \sum_{i=1}^{n-1} \sum_{j=i+1}^{n} r_{i-1} r_{j-i-1} r_{n-j}$ for $n > 1$, it holds that $q_n \geq r_n$ for all n.

Proof: The proof is complicated and not especially enlightening, and Theorem 5 will allow a computer to quickly prove the proposition (the last paragraph of this proof, which is simple and straightforward is still necessary). This proof is included to illustrate the headache that Theorem 5 will help alleviate. The main step of the proof is to show that Q_n contains $R_{n,T}$ where $T = (45312, 10101)$. We will show that every permutation in $R_{n,T}$ avoids 2143, 2413, and 3142. First, note that for 2413, and 3142, no proper subword with length greater than 2 contains only consecutive numbers. When we divide a permutation into subwords to fit into the template, each subword must contain only consecutive numbers. Thus we can conclude that if a pattern is present in a permutation in $R_{n,T}$, then it is contained entirely in a single subword or each element is in a different subword. The second case cannot occur because the permutation 45312 avoids both patterns. To see that the first case cannot occur, suppose by way of contradiction that it does, and pick n minimally so that a permutation of $R_{n,T}$ contains one of the two patterns under consideration. When we divide up this permutation into subwords so that it fits into T, we must choose some subword to contain pattern, but then this subword is a shorter permutation which contains the pattern, providing a contradiction. This shows that every permutation in $R_{n,T}$ avoids 2413 and 3142.

Next we will see that every permutation also avoids 2143. Again suppose by way of contradiction that there is a permutation in $R_{n,T}$ which contains 2143, and pick n minimally so that this occurs.

Then, if we divide up the permutation into 5 subwords, $W_1, ..., W_5$, which fit the template T the occurrence of 2143 cannot be contained entirely in any one subword. Therefore, W_1 either contains no part of the occurrence, contains the 2, or contains the 21. In the first two cases W_2 must not contain any part of the occurrence either; it cannot contain the 2 or 1 because it is the largest element of the permutation. If W_1 was empty, then we must fit 2143 into $W_3W_4W_5$, which is impossible because either the 2 will go in W_3 even though each element of W_3 must be greater than each element of W_4 and W_5, or else we would need to fit 143 into W_5 which can't happen because they are not consecutive integers (the 2 is missing). If W_1 contained 2, then W_2 must contain 4 and 3 because all the elements of every other subword must be less than the elements of W_1. Therefore, the permutations in $R_{n,T}$ avoid 2143, and so $R_{n,T} \subseteq Q_n$.

Now, we just need to show that $|R_{n,T}| = r_n$. First, it follows from the definition of $R_{n,T}$ that $|R_{0,T}| = |R_{1,T}| = 1$. Then, for a permutation in $R_{n,T}$, we will say that n occurs at position i and 1 at position j. We get that $1 \leq i \leq n-1$ and $i+1 \leq j \leq n$. Then, W_1 can be any of the r_{i-1} elements of $R_{i-1,T}$, W_3 can be any of the r_{j-i-1} elements of $R_{j-i-1,T}$, and W_5 can be any of the r_{n-j} elements of $R_{n-j,T}$. Therefore, $r_n = \sum_{i=1}^{n-1} \sum_{j=i+1}^{n} r_{i-1} r_{j-i-1} r_{n-j}$ for $n > 1$.

While this recurrence for (r_n) is reminiscent of the Catalan recurrence, it does not appear to have a similarly nice closed form solution.

Fortunately, it is possible to prove results of this kind experimentally without the need for detailed write-ups. The following theorem establishes a sufficient condition for $R_{n,T}$ to avoid a set of patterns which is independent of n, and so can be tested for all n at once using a computer.

Theorem 5: Let $T = (P, B)$ be a template, let B have k 0's, and let σ be a pattern of length $l > 0$. Then, if there exists an n such that $R_{n,T}$ contains a permutation which has σ as a pattern, there also exists an $m \leq (l-1)(k+1)+1$ such that $R_{m,T}$ also contains a permutation which has σ as a pattern.

Proof: This Theorem will be an immediate corollary of Theorem 6, and, while it can be proved separately, the proof is almost identical to that of Theorem 6, so we omit it.

With this result in hand, the author's laptop was able to prove Proposition 4 in 16 seconds using Maple.

There is no particular reason to consider templates just one at a time. Analogously to how we originally defined templates, we define the set of length n permutations corresponding to the set of templates $\mathcal{T} = \{T_1, ..., T_r\}$. We will call this set of permutations $S_{n,\mathcal{T}}$, and define it recursively as follows. First, $S_{0,\mathcal{T}}$ is the empty string and $S_{1,\mathcal{T}} = \{1\}$ regardless of $\mathcal{T}$. Then, $S_{n,\mathcal{T}}$ is the set of permutations π of length n such that, for some $T = (P, B) \in \mathcal{T}$, we can divide π into subwords $W_1, ..., W_t$ such that if $p_i > p_j$, then every element of W_i greater than every element of W_j. Moreover, we require that each W_i of length l be an element of $S_{l,\mathcal{T}}$ (rather than of $R_{l,T}$), and, if $B_i = 0$, then W_i must have exactly one element.

We will finish this section by proving a generalization of Theorem 5 for sets of templates, and giving an example of its application.

Theorem 6: Let $\mathcal{T} = \{(P_1, B_1), ..., (P_r, B_r)\}$ be a set of templates, suppose that for all i B_i has no more than k 0's, and let σ be a pattern of length $l > 0$. Then, if there exists an n such that $S_{n,\mathcal{T}}$ contains a permutation which has σ as a pattern, there also exists an $m \leq (l-1)(k+1)+1$ such that $S_{m,\mathcal{T}}$ also contains a permutation which has σ as a pattern.

Proof: Fix k and n; we proceed by induction on l. If $l = 1$, then σ is the pattern 1 and is contained in the permutation 1 which is the element of $S_{1,\mathcal{T}}$. Assume that the theorem holds for patterns of length up to $l-1$. Now let $\pi' \in S_{n,\mathcal{T}}$ be the permutation which contains σ as a pattern, and pick some occurrence of σ in π'. We can choose $T = (P, B) \in \mathcal{T}$ and divide π' into subwords $W'_1, ..., W'_t$ (where $t = |P|$) such that the W'_i fit the template T. We can similarly divide σ into subwords $U_1, ..., U_t$ so that U_i is the portion of the chosen occurrence of σ which lies in W'_i. If there exists i such that only U_i is nonempty, then W'_i contains σ and is shorter than π', so set $\pi' = W'_i$ and repeat the decomposition for the new π'. Repeat until either at least two U_i are nonempty or $|\pi'| \leq (l-1)(k+1)+1$. In the second case we are done, so assume that the first case holds.

We will now find m and construct a permutation $\pi \in S_{m,\mathcal{T}}$ which contains σ. Like π' we need to be able to divide π into $W_1, ..., W_t$ to fit T, so we will construct the W_i individually. For each i, let $u_i = |U_i|$. By the induction hypothesis, there exist W_i such that $|W_i| = w_i \leq (u_i - 1)(k+1)+1$, $W_i \in S_{w_i,\mathcal{T}}$, and U_i is a pattern in W_i. It may be that for some i, W_i is empty even though $B_i = 0$; if this is the case, we must add up to k new W_i to ensure that each W_i has length 1 whenever $B_i = 0$. Lastly, we choose i so that $p_i = 1$ and j so that $p_j = 2$ and increase every element of W_j by the same amount so that every element of W_j is greater than every element of W_i, and we repeat this with $j = 3..t$ and $i = j-1$. Now, concatenating all the W_i gives a permutation π of length m in $S_{m,\mathcal{T}}$ which contains the pattern σ.

It remains to show that $m \leq (l-1)(k+1)+1$. Let $I = \{i : u_i > 0\}$; using the construction of π and the induction hypothesis, we find that $m \leq \sum_{i \in I}((u_i - 1)(k+1)+1) + k = (k+1)(\sum_{i \in I} u_i) - k \cdot |I| + k = (k+1)(l) - k|I| + k \leq (k+1)(l-1)+1$ because we found at the end of the first paragraph that $|I| \geq 2$. Therefore, the proof is complete by induction.

Theorem 6 can give lower bounds on the sizes of many sets of avoidance classes. As an example, we offer the following proposition:

Proposition 7: Let Q_n be the set of all permutations of n which avoid every element of $\{2341, 2413, 2431, 3241\}$ and let $q_n = |Q_n|$. Then, if the sequence $(s_n)_{n=0}^{\infty}$ is defined by $s_0 = s_1 = 1$, and $s_n = \sum_{i=1}^{n-1} \sum_{j=i+1}^{n} 2 \cdot s_{i-1} s_{j-i-1} s_{n-j}$ for $n > 1$, it holds that $q_n \geq s_n$ for all n.

Proof: Let $T_1 = (14253, 10101), T_2 = (15243, 10101)$, and $\mathcal{T} = \{T_1, T_2\}$. Using Maple, one can generate $S_{n,\mathcal{T}}$ for $1 \leq n \leq 10$, and confirm that every permutation in each of these sets avoids 2341, 2413, 2431, and 3241 (we did this on a laptop in less than 7 minutes). Because these patterns all have length 4, both B_1 and B_2 have two 0's, and $(4-1) \cdot (2+1) + 1 = 10$, Theorem 6 promises that, for all n, every permutation in $S_{n,\mathcal{T}}$ avoids 2341, 2413, 2431, and 3241.

Now we show that $|S_{n,\mathcal{T}}| = s_n$ by induction. Certainly $|S_{0,\mathcal{T}}| = |S_{1,\mathcal{T}}| = 1$. When picking a permutation in $S_{n,\mathcal{T}}$, we first choose whether this permutation will follow the template T_1 or T_2. This will not cause us to count any permutation twice because if a permutation has $n-1$ appear before n, then it can only follow T_1, and if it has n appear before $n-1$ then it can only follow T_2. Now, for T_1, we must choose the location of $n-1$, call this i, and the location of n, call it j. For T_2, we will call the location of n i and the location of $n-1$ j. For either template, we have $1 \leq i \leq n-1$, $i+1 \leq j \leq n$. Once i and j are chosen, we can fill in the portion of the permutation before position i in any of s_{i-1} ways, the portion between positions i and j in s_{j-i-1} ways, and the portion following position j in s_{n-j} ways. Therefore, $|S_{n,\mathcal{T}}| = \sum_{i=1}^{n-1} \sum_{j=i+1}^{n} 2 \cdot s_{i-1} s_{j-i-1} s_{n-j} = s_n$ for $n > 1$.