

Pattern Popularity in 132-avoiding Permutations

Kate Rudolph*

Department of Mathematics
Massachusetts Institute of Technology
Cambridge, MA, U.S.A.

k8r@mit.edu

Submitted: Aug 13, 2012; Accepted: Jan 3, 2012; Published: Jan 14, 2013

Mathematics Subject Classifications: 05A05, 05A15, 05A19

Abstract

The popularity of a pattern p is the total number of copies of p within all permutations of a set. We address popularity in the set of 132-avoiding permutations. Bóna showed that in this set, all other non-monotone length-3 patterns are equipopular, and proved equipopularity relations between some length- k patterns of a specific form. We prove equipopularity relations between general length- k patterns, based on the structure of their corresponding binary plane trees. Our result explains all equipopularity relations for patterns of length up to 7, and we conjecture that it provides a complete classification of equipopularity in 132-avoiding permutations.

1 Introduction

Let $\sigma = \sigma_1 \dots \sigma_n$ be a permutation of the numbers $1, \dots, n$, and $p = p_1 \dots p_k$ be a permutation of $1, \dots, k$, with $k \leq n$. The permutation σ *contains* the pattern p if some subsequence of σ is order-isomorphic to p ; that is, if there is some sequence of indices $1 \leq i_1 < \dots < i_k \leq n$ such that $\sigma_{i_s} < \sigma_{i_t}$ if and only if $p_s < p_t$. A permutation may contain multiple copies of a pattern: for example, the permutation 51243 contains two copies of the pattern 123, one at indices 2, 3 and 4, and another at indices 2, 3, and 5. If σ contains no copies of a pattern p , we say that σ *avoids* p . The set of all $\sigma \in S_n$ which avoid p is denoted $S_n(p)$.

Let $f(\sigma, p)$ be the number of copies of p occurring in σ .

Definition 1. The *popularity* $P_S(p)$ of a pattern p in a set S of permutations is given by:

$$P_S(p) = \sum_{\sigma \in S} f(\sigma, p).$$

*Supported by NSF grant DMS-1062709 and NSA grant H98230-11-1-0224

Example 2. If we take $S = S_n$, the set of all permutations of $1, \dots, n$, then any two patterns of the same length are equally popular. By linearity of expectation, for all patterns p of length k , $P_{S_n}(p) = \frac{n!}{k!} \binom{n}{k}$.

A notion related to popularity was defined by Bóna, Sagan, and Vatter in [5]. They define the *frequency* $S_{n,q}(c)$ to be the number of permutations of length n which contain exactly c copies of the pattern q . The popularity of q in S_n is related to the frequency by $P_{S_n}(q) = \sum_{c \geq 0} c * S_{n,q}(c)$. That is, the popularity is the total number of occurrences of the pattern q , counting multiplicity. Frequency is not studied in this paper, but we mention it to clarify the new term popularity.

While our example of popularity was easy, it becomes more difficult to compute the popularity of patterns when $S = S_n(p)$ for some pattern p . This problem was first proposed, in slightly different terms, by Joshua Cooper [3], who asks: given two permutations p and r , what is the expected number of copies of p in a permutation chosen at random from $S_n(r)$? Instead of considering expectation, in this paper we will examine the equivalent question of how many copies of p exist in all permutations in $S_n(r)$, which is exactly the popularity.

The focus of this paper is the popularity of patterns in $S_n(132)$, the set of length- n 132-avoiding permutations. For convenience, we will write $P_{S_n(132)}(p)$ as $A_n(p)$.

Permutations of length n avoiding 132 are counted by the Catalan numbers. The popularity of any pattern p containing 132 in $S_n(132)$ will be zero, because if $\sigma \in S_n(132)$ contains p , then σ contains 132, a contradiction. Thus, we will look only at the popularity of patterns which are themselves permutations in $S_n(132)$.

Bóna [1] showed that for all patterns p of length k and for all n , we have

$$A_n(12 \cdots k) \leq A_n(p) \leq A_n(k \cdots 21);$$

that is, the strictly increasing pattern is the least popular and the strictly decreasing pattern is the most popular among all patterns of length k .

If $A_n(p) = A_n(q)$ for all n , we say that p and q are *equipopular*. This is an equivalence relation, so we can divide patterns of length k into equivalence classes using equipopularity. Bóna [2] showed that for length-3 patterns, there are three equivalence classes: 123 and 321 comprise the least-popular and most-popular classes, respectively, while 213, 312, and 231 are all equipopular.

In [2], Bóna also proved that patterns of a certain specific form are equipopular. Adopting the notation in Definition 4, Bóna's theorem states

Theorem 3 (Bóna). *Let q and t be any non-empty patterns that end in their largest entry, and let i_u denote the increasing pattern $12 \dots u$. Then for all positive integers n , we have*

$$A_n((q \ominus t) \oplus i_u) = A_n((q \oplus i_u) \ominus t).$$

This theorem provides an exponential number of pairs of patterns which are equipopular, but at least one pattern in the pair must end in its largest element in order for the

theorem to apply. Thus if a pattern p is not equipopular to any pattern q ending in its largest element, then Bóna's theorem of [2] can not apply to p .

This paper extends Bóna's results and proves that more general classes of patterns are equipopular. Our results completely explain all experimentally observed equipopularity classes in patterns of length at most 7. The main theorem uses a bijection from 132-avoiding permutations to another combinatorial object called binary plane trees, and states that when two binary plane trees have the same spine structure (see Definition 10), their corresponding patterns are equipopular. Thus the result applies to all patterns of all lengths k . We have experimentally computed the popularities for all patterns with $k \leq 7$, and there are no additional equipopularity relations besides those predicted by our theorem. We conjecture that our theorem also provides a full classification of the equipopularity classes for all k .

For example, the patterns 43251, 43521, 53214, 53241, and 54213 are all equipopular by the main theorem of this paper. Calculations of their popularities $A_n(p)$ for the first few values of n show that no other length-5 patterns are equivalent to these five. Since none of them ends in its largest element, they cannot be proved equivalent by Bóna's theorem of [2]. In general, the main result of this paper explains all observed equipopularity results in our numerical evidence, a problem left open by Bóna in [2].

In Section 2, we use the method of generating functions outlined by Bóna in [1] to prove general results about extending equipopular patterns, which will be used in the proof of the main theorem in Section 4. In Section 3 we describe a bijection from 132-avoiding permutations to binary plane trees, and define the notation which allows us to precisely state the main theorem. Section 4 contains the proof of the main theorem.

2 Arguments Using Generating Functions

In this section we will show that, given equipopular patterns, we can build longer patterns from them which are also equipopular. We will do this by manipulating the generating functions for the popularity. The theorems proved here will be used in the proof of the main theorem.

Bóna [1] lays out a procedure for breaking down the possible cases of the occurrence of a pattern p in a 132-avoiding permutation σ of length n .

In any 132-avoiding permutation σ , all the entries to the left of the entry n must be greater than all the entries to the right of n , otherwise there would be an occurrence of the pattern 132. We say that the pattern p has a *cut* at index i if all entries of p before and at index i are greater than all entries after index i . The pattern p can occur in σ in several ways, enumerated by Bóna [1]:

1. p occurs entirely to the left of n
2. p occurs entirely to the right of n
3. the largest element k in p can be represented by the n in σ , with the rest of the pattern p split between the left and right sides of σ

4. for each cut in p , we can split p between the left and right blocks of σ at that cut.

If p ends in its largest element, the fourth case cannot occur as there are no such indices i .

Let c_n be the n th Catalan number, and $C(x) = \sum_{n \geq 0} c_n x^n$ be the generating function for the Catalan numbers.

We will show that given two equipopular patterns p and q , it is possible to construct more equipopular patterns of a longer length. This will require notation for combining two patterns to make another. We define two operations for combining patterns.

Definition 4. Given two patterns p and q , of lengths k and m respectively, we let $p \oplus q$ be the pattern of length $k + m$ consisting of p followed by q shifted up by k :

$$(p \oplus q)_i = \begin{cases} p_i & \text{if } i \leq k \\ q_{i-k} + k & \text{if } i > k. \end{cases}$$

Similarly, $p \ominus q$ is the pattern of length $k + m$ containing p shifted up by m , followed by q :

$$(p \ominus q)_i = \begin{cases} p_i + m & \text{if } i \leq k \\ q_i & \text{if } i > k. \end{cases}$$

Example 5. If $p = 3421$ and $q = 123$, then $p \oplus q = 3421567$ and $p \ominus q = 6754123$.

Theorem 6. 132-avoiding patterns p and q are equipopular if and only if $p \oplus 1$ and $q \oplus 1$ are equipopular.

We note that if a pattern p avoids 132, then so must $p \oplus 1$, because the additional 1 at the end cannot introduce an occurrence of 132.

Proof. The idea of this proof is to relate the generating function for the popularity of p to that of $p \oplus 1$.

Let $a_n = A_n(p)$ and $b_n = A_n(q)$. Let $A(x) = \sum_{n \geq 0} a_n x^n$ be the generating function for a_n and $B(x) = \sum_{n \geq 0} b_n x^n$ be the generating function for b_n . Let $a'_n = A_n(p \oplus 1)$ and $b'_n = A_n(q \oplus 1)$.

We will proceed by counting the number of ways $p \oplus 1$ can occur in a permutation whose entry n is at index i , and then we will sum over all i .

Since $p \oplus 1$ ends in its largest element, there are no cuts, so we only need to consider the first three cases.

1. there are a'_{i-1} ways for $p \oplus 1$ to occur to the left of the n , and for each of those there could be any 132-avoiding permutation on the right of the n , for c_{n-i} possibilities in each case ($a'_{i-1} c_{n-i}$ possibilities).
2. there are a'_{n-i} ways for $p \oplus 1$ to occur on the right of the n and for each of those, c_{i-1} ways to fill the elements on the left side of the n ($a'_{n-i} c_{i-1}$ possibilities).
3. if the largest element of $p \oplus 1$ coincides with n , an instance of p must occur to the left of n and any 132-avoiding permutation may fall to the right ($a_{i-1} c_{n-1}$ possibilities).

Summing over all possible positions of i we get

$$a'_n = \sum_{i=1}^n a'_{i-1} c_{n-i} + \sum_{i=1}^n a'_{n-i} c_{i-1} + \sum_{i=1}^n a_{i-1} c_{n-i}$$

Converting to a generating function by taking $A'(x) = \sum_{n \geq 0} a'_n x^n$, we find

$$A'(x) = 2xA'(x)C(x) + xA(x)C(x)$$

which is equivalent to

$$A'(x) = \frac{x A(x) C(x)}{1 - 2x C(x)}.$$

The same argument shows that

$$B'(x) = \sum_{n \geq 0} b'_n x^n = \frac{x B(x) C(x)}{1 - 2x C(x)}.$$

Thus $A(x) = B(x)$ if and only if $A'(x) = B'(x)$, so p and q are equipopular if and only if $p \oplus 1$ and $q \oplus 1$ are equipopular. $\square$

If we let I_j be the pattern $123 \dots j$, then repeated application of Theorem 6 gives the following corollary.

Corollary 7. *132-avoiding patterns p and q are equipopular if and only if $p \oplus I_j$ and $q \oplus I_j$ are equipopular for any positive integer j .*

If p and q are equipopular, $p \ominus 1$ and $q \ominus 1$ need not be equipopular. For example, although 231 and 213 are equipopular [2], we have $231 \ominus 1 = 3421$ and $213 \ominus 1 = 3241$. A little computation shows $A_5(3241) = 15$ and $A_5(3421) = 16$, so 3421 and 3241 are not equipopular. However, when p and q both end in their largest element, we have the following result:

Theorem 8. *132-avoiding patterns p and q are equipopular if and only if $(p \oplus 1) \ominus 1$ and $(q \oplus 1) \ominus 1$ are equipopular.*

Proof. We will find a relation between the generating function for p and that for $(p \oplus 1) \ominus 1$. As before, let $a_n = A_n(p)$, $b_n = A_n(q)$, $a'_n = A_n(p \oplus 1)$ and $b'_n = A_n(q \oplus 1)$. Also, let $a_n^* = A_n((p \oplus 1) \ominus 1)$ and let $b_n^* = A_n((q \oplus 1) \ominus 1)$. Let $A(x)$, $B(x)$, $A'(x)$, $B'(x)$, $A^*(x)$, and $B^*(x)$ be the generating functions for a_n , b_n , a'_n , b'_n , a_n^* , and b_n^* , respectively. From the proof of Theorem 6, we know

$$A'(x) = \frac{x A(x) C(x)}{1 - 2x C(x)} \quad \text{and} \quad B'(x) = \frac{x B(x) C(x)}{1 - 2x C(x)}.$$

Now, we will use the cases outlined at the beginning of this section to count the number of ways $(p \oplus 1) \ominus 1$ can occur in a length- n 132-avoiding permutation whose entry n is at index i . Later we will sum over all i . The four cases, in the same order as before, are:

1. the $(p \oplus 1) \ominus 1$ pattern occurs entirely to the left of index i and the right of index i is occupied by any 132-avoiding permutation $(a_{i-1}^* c_{n-i})$.
2. the $(p \oplus 1) \ominus 1$ pattern occurs entirely to the right of index i and the left of index i is occupied by any 132-avoiding permutation $(a_{n-i}^* c_{i-1})$.
3. the largest element in the $(p \oplus 1) \ominus 1$, which is the second to last element in the pattern, coincides with the n in σ at index i . A p pattern occurs to the left. On the right of the n is any 132-avoiding permutation, but we must also choose which of the $n - i$ elements to the right will be included to form the last element of our $(p \oplus 1) \ominus 1$ pattern. $(a_{i-1} c_{n-1}(n - i))$.
4. There is only one cut: the $(p \oplus 1)$ pattern occurs on the left side of the n and the last element of the $(p \oplus 1) \ominus 1$ pattern occurs somewhere on the right side of the n $(a'_{i-1} c_{n-i}(n - i))$.

Summing over all i , we find that

$$a_n^* = \sum_{i=1}^n a_{i-1}^* c_{n-i} + \sum_{i=1}^n a_{n-i}^* c_{i-1} + \sum_{i=1}^n a_{i-1} c_{n-i}(n - i) + \sum_{i=1}^n a'_{i-1} c_{n-i}(n - i)$$

Let $Z(x) = \sum_{n \geq 0} n c_n x^n$. Then we find

$$A^*(x) = 2xA^*(x)C(x) + xA(x)Z(x) + xA'(x)Z(x),$$

or equivalently

$$A^*(x) = \frac{xZ(x)(A(x) + A'(x))}{1 - 2xC(x)} = \frac{xZ(x) \left(A(x) + A(x) \frac{xC(x)}{1 - 2xC(x)} \right)}{1 - 2xC(x)}.$$

The same argument shows that

$$B^*(x) = \frac{xZ(x)(B(x) + B'(x))}{1 - 2xC(x)} = \frac{xZ(x) \left(B(x) + B(x) \frac{xC(x)}{1 - 2xC(x)} \right)}{1 - 2xC(x)}.$$

Thus $A^*(x) = B^*(x)$ if and only if $A(x) = B(x)$, so p and q are equipopular if and only if $(p \oplus 1) \ominus 1$ and $(q \oplus 1) \ominus 1$ are equipopular. $\square$

3 Bijection to Binary Plane Trees

In order to prove the main theorem, we introduce a set of objects, binary plane trees, which like 132-avoiding permutations can be counted by the Catalan numbers.

Definition 9. A *binary plane tree* is a rooted, unlabeled tree in which each node has at most two children, and every child is designated as either the left or the right child of its parent, even if it is the only child.




Figure 1: An example of the binary plane tree $T(5647213)$. The tree is labeled to illustrate the bijection, although in general binary plane trees are unlabeled.

There is a well-known bijection, described in [2], between 132-avoiding permutations and binary plane trees. For a 132-avoiding permutation p , we construct its associated binary plane tree $T(p)$ recursively. If p is the empty sequence, $T(p)$ is the empty tree. Then suppose p is of length n . The root node of $T(p)$ represents the entry n of p . If p_1 is the subsequence of p before the n entry, and p_2 is the subsequence of p after the n entry, then in $T(p)$ the left subtree of the root is $T(p_1)$ and the right subtree is $T(p_2)$.

We can recover the 132-avoiding permutation p corresponding to a tree by labeling each of the nodes by the entry it must represent. (Although binary plane trees are unlabeled, the labeling we give them is uniquely determined by the shape of the tree, so it does not affect the combinatorial definition of binary plane trees.) We know the root node must represent the element n . Suppose the right subtree has i nodes, and thus the left subtree has $n - i - 1$ nodes. Since every element to the left of n in p is greater than every element to the right of n (because otherwise there would be an instance of the pattern 132) we know that the left subtree of the root must represent the values $i + 1, \dots, n - 1$ and the right subtree must represent the values $1, 2, \dots, i$. Recursively applying this rule allows us to construct a unique labeling of the nodes on the tree. The permutation it represents can be recovered by an in-order reading of the nodes: for each node we encounter, perform the in-order reading of its left child, then read the value of the node itself, then perform the in-order reading of its right child.

Definition 10. The *spines* of a tree $T(p)$ are the connected components of $T(p)$ when all edges connecting left children to their parents are deleted. The *length* of a spine is the number of nodes in the spine. The *spine structure* of $T(p)$ is the sequence of the lengths of spines in $T(p)$, sorted in descending order.

An example of the spine structure is given in Figure 2. Note that the spine structure of a tree $T(p)$ with n nodes is just a partition of n . Now we can precisely state the main theorem of this paper.

Theorem 11. *For 132-avoiding patterns p and q , if the spine structures of $T(p)$ and $T(q)$ are the same, then p and q are equipopular.*




Figure 2: An example of a binary plane tree with spine structure $\{3, 2, 1, 1\}$. On the right the spines are highlighted.

4 Proof of Main Theorem

Several of the ideas in this proof are borrowed from [1]: he introduces the idea of looking for subtrees in binary plane trees as equivalent to looking for patterns in permutations, and the idea of switching some subtrees.

To begin the proof, we first notice that switching two left subtrees in a binary plane tree does not affect its spine structure. If our theorem is true the patterns represented by the trees before and after the switch should have the same popularity. This motivates our left-subtree-switching lemma.

Lemma 12 (Left-Subtree-Switching Lemma). *Given a 132-avoiding pattern p of length k , and two nodes a and b of $T(p)$ such that neither is in the left subtree of the other, let $T(q)$ be the tree obtained by switching the left subtrees of a and b , and q be the pattern represented by this tree. Then p and q are equipopular.*

We allow the case where a or b has an empty left subtree. Figure 3 illustrates a left subtree switch.




Figure 3: On the left, $T(5647213)$, and on the right, $T(5463712)$. The right subtree is obtained from the left by switching the left subtrees of the two enlarged nodes. Lemma 12 states that 4657213 and 5463712 are equipopular.

Proof. Let B_n be the set of all n -node binary plane trees $T(\sigma)$, k of whose nodes are colored red such that those k nodes represent a p pattern in σ . The order of B_n is just the popularity of p , i.e. $|B_n| = A_n(p)$. Similarly, let C_n be the set of all n -node binary trees

$T(\sigma)$, k of whose nodes are colored red, such that those k nodes represent a q pattern in σ . So $|C_n| = A_n(q)$. We will prove the lemma by constructing a bijection φ from B_n to C_n .

Let i be the index of a in the in-order reading of $T(p)$ and let j be the index of b . Then given some element Q of B_n , let Q_a be the i th red node in the in-order reading of Q and Q_b be the j th red node. Without loss of generality, let $i < j$. Then $\varphi(Q)$ is the tree obtained by switching the left subtrees of Q_a and Q_b , without changing whether any node is colored red. We must show that this is well defined and that it is a bijection between B_n and C_n .

To show φ is well defined, we must show that Q_a is not in the left subtree of Q_b and vice versa. Since $i < j$, a is before b in the in-order reading of $T(p)$, and Q_a is before Q_b in Q . Thus Q_b is not in the left subtree of Q_a . Since a comes before b in $T(p)$ but is not in the left subtree of b , because of the way binary plane trees are labeled to recover a 132-avoiding permutation we can conclude that the label of a is greater than the label of b . So the i th element of p is greater than the j th element of p . Thus, since the red nodes in Q form a p pattern, the label of the i th red node in Q , or Q_a , must be greater than the label of the j th red node, Q_b . If Q_a were in the left subtree of Q_b , its label would be less than Q_b , so Q_a must not be in the left subtree of Q_b . Thus function φ is well defined on elements of B_n .

Now we will show that for $Q \in B_n$, we always have $\varphi(Q) \in C_n$, that is, the red nodes in $\varphi(Q)$ form a q pattern.

When we switch from $T(p)$ to $T(q)$, the only nodes that move are the nodes in the left subtree of a , which we will call α , and the nodes in the left subtree of b , which we will call β . In $T(q)$, the sets of nodes α and β simply switch position, and nothing else changes its relative order.

So, let $Q \in B_n$. The red nodes in Q represent a p pattern, so all the red nodes corresponding to the α nodes in $T(p)$ must be in the left subtree of Q_a , since they must be before Q_a and have label less than Q_a . No red node in the left subtree of Q_a can correspond to an element of the pattern p that is not in α since α consists of all the pattern elements less than a and to the left of a . So the red nodes in the left subtree of Q_a are exactly those representing α . Similarly, the red nodes in the left subtree of Q_b are exactly those representing β . When we switch the left subtrees of Q_a and Q_b , the only effect on the red nodes is to switch the positions of the nodes representing α and β . Thus, we've produced a q pattern.

Now, to show that φ is a bijection, we will show that its inverse is a very similar function φ' . Recall that i and j are the indices of the nodes a and b in the in-order reading of $T(p)$. Let i' be the index of a in the in-order reading of $T(q)$. Then $i' = i - |\alpha| + |\beta|$, since i' is the index when we take out all of the elements of α that appear before i and replace them with all the elements of β . The index of b in $T(q)$ is the same as the index of b in $T(p)$, j , because the only things that are moved around all occur before b in the in-order reading, and there are the same number of them before and after the switch, so b still has index j .

Then, we can see that to invert φ , which switches the left subtrees of the i th and j th




Figure 4: Though the two trees have the same spine structure $\{3, 2, 1\}$, neither can be transformed to the other by left-subtree switches.

red nodes, we must simply switch those left subtrees back. But the index i has changed to i' since the subtrees we switched may not have been the same size. So the inverse function $\varphi' : C_n \rightarrow B_n$ is defined as follows. Let $U \in C_n$, and let $U_{i'}$ be the i' 'th red node in the in-order reading of U , and U_j the j 'th red node. Then $\varphi'(U)$ is the tree obtained by switching the left subtrees of $U_{i'}$ and U_j .

We can see that φ' is the inverse of φ because it simply switches the subtrees back to their original positions.

Thus φ is a bijection from B_n to C_n , so $|B_n| = |C_n|$ and $A_n(p) = A_n(q)$. Since this bijection can be defined for arbitrary n , $A_n(p) = A_n(q)$ for all n and so p and q are equally popular. $\square$

Thus, every tree that can be obtained by repeatedly switching left subtrees of $T(p)$ represents a pattern equipopular to p . It would be convenient if every tree with the same spine structure could be transformed, via left-subtree switches, into any other tree with that spine structure; however, this is not the case. For example, Figure 4 shows two trees which cannot be transformed into each other by any sequence of left subtree switches.

However, we can show that every tree can be transformed to a “left-justified” tree with the same spine structure via left subtree switches.

Definition 13. A *left-justified* binary plane tree is one in which every node that is a right child of its parent does not have a left child.

Thus the only nodes in a left-justified tree which are left children must have as parents either nodes which are also left children, or else the root. So in a left-justified tree, all the down-left edges, between a parent and its left child, form a path down the left side of the tree. We call this path the *trunk* of the tree. The right subtree of each node along the trunk forms a spine. For each spine, we call the node in that spine that is on the trunk the *top* node of that spine. Similarly, the spine whose top is at the root is the *topmost* spine, and the *lowest* spine in a left-justified tree is the one whose top node occurs lowest on the trunk.

Lemma 14. For each binary plane tree $T(p)$, we can construct a left-justified tree $T^*(p)$ via only left-subtree switches.

Proof. If $T(p)$ is already left-justified, we are done. Otherwise, let S be the set of nodes in $T(p)$ which are right children and have a left child. We will describe an operation on $T(p)$ which will reduce the size of S , and so repeated application of this operation will reduce the size of S to zero, meaning that the resulting tree will be left-justified.

Start at the root of $T(p)$, and move to the left child of the current node until we reach a node x which has no left child. Let y be some element of S . We know that y is not in the left subtree of x because x has no left child. Also, x is not in the left subtree of y because the only ancestors of x are themselves left children, and y is in S so it must be a right child. So we may use Lemma 12 and switch the left subtrees of x and y . Since x has no left child, after the switch y has no left child, so y is no longer in S . Every other node that was in S is still in S because we haven't moved any of their left children. No new nodes are added to S because the only node which gained a left child was x , but x is not a right child of its parent so it cannot be an element of S .

Repeat this process until S is empty. The resulting tree is left-justified. $\square$

Switching left subtrees preserves popularity of the patterns represented by the trees, and it also preserves the spine structure of the trees. So this shows us that that any pattern p is equipopular to some pattern q such that $T(p)$ and $T(q)$ have the same spine structure, but $T(q)$ is left-justified. However, this is still not enough to prove our theorem: two left-justified trees may have the same spine structure, but the spines may appear in a different order. We need a way to permute the order of the spines along the trunk in a left-justified tree while preserving the popularity of the pattern that the tree represents.

Lemma 15. *If $T(p)$ and $T(q)$ are left-justified trees with the same spine structure, and the top spines of $T(p)$ and $T(q)$ have the same length $\ell \geq 2$, then p and q are equipopular.*

Proof. We will prove this lemma by showing that every left-justified tree $T(p)$ whose top spine has length greater than 2 can be transformed via a series of left subtree switches to a left-justified tree with all spines except the top spine appearing in sorted order by length along the trunk.

We present an algorithm for sorting all of the spines except the top spine along the trunk. Begin by letting a be the root node of $T(p)$ and b be the right child of a . We will use b as a sort of holding cell: at the beginning, b cannot have a left child because $T(p)$ is left-justified, so any left subtree switch performed with b as one of the elements being switched just moves a left subtree to b .

The algorithm is performed by repeating the following process until the spines are in sorted order from shortest length to longest, with the longer spines further from the root, except the spine starting from the root.

1. Switch the left subtree of a with the left subtree of b .
2. Determine the spine with the shortest length which is in the left subtree of b . Let x be the top node of this spine.
3. Switch the left subtree of x 's parent with the left subtree of a .

4. Let y be the top node of the lowest spine in the left subtree of a .
5. Switch the left subtree of b with the left subtree of y .
6. Redefine a to be the node currently labeled x .

We claim that after r repetitions of these steps, the r shortest spines of $T(p)$ occur immediately after the topmost spine, in sorted order, and that a refers to the top of the last spine of these r sorted spines. We proceed by induction.

After 0 repetitions, our claim is vacuously true. Suppose that the claim is true for $r = k$, so we have repeated the steps k times and the smallest k spines besides the top spine are in sorted order along the trunk. The inductive hypothesis tells us that a is the top node of the lowest of these k spines. So the entire left subtree of a consists of the unsorted spines. We carry out the loop for the $(k + 1)$ st time. In step 1, the left subtree of a is moved to b , which does not have a subtree. Thus we've moved all the unsorted spines into the left subtree of b . In step 2, we determine x , the top node of the shortest spine in this subtree. But the shortest spine of all the unsorted spines is the one which needs to go next. In step 3, we switch the left subtree of x 's parent with the left subtree of a . Since we had already moved a 's left subtree away, it does not have a left subtree, so this is equivalent to moving the left subtree of x 's parent to the left subtree of a , and thus the spine whose top is x is now next in the sorted order, as desired. Steps 4 and 5 move the left subtree of b back to the end of the trunk, left-justifying the tree again. Finally, in step 6 we redefine a to be one node lower on the trunk. Thus after the loop we have sorted $k + 1$ spines, since we added the spine whose top is x to the sorted order, and a is the top of the lowest of those $k + 1$ sorted spines. Our induction is complete.

Therefore, repeating the loop a number of times equal to the number of spines will ensure that all the spines are sorted.

To prove the lemma, sort $T(p)$ and $T(q)$ according to the above process, to obtain the trees $T(p')$ and $T(q')$. Since $T(p)$ and $T(q)$ both have the same spine structure and the same length of the topmost spine, then $T(p')$ and $T(q')$ are identical, so the patterns p' and q' they represent are the same. Since we performed only left subtree switches, p and p' are equipopular and q and q' are equipopular by the Left-Subtree-Switching Lemma. But p' and q' are the same so p and q are equipopular. $\square$

We are ready to prove our main theorem, restated here for convenience.

Theorem 16. *For 132-avoiding patterns p and q , if the spine structures of $T(p)$ and $T(q)$ are the same, then p and q are equipopular.*

Proof of Main Theorem. We are given p and q such that $T(p)$ and $T(q)$ have the same spine structure. Using Lemma 14, we transform $T(p)$ to the left-justified tree $T(p_1)$. Since we used only left subtree switches, by Lemma 12 we know p and p_1 are equipopular. Because left subtree switches preserve spine structure, $T(p)$ and $T(p_1)$ have the same spine structure. Similarly, we construct $T(q_1)$, so that q and q_1 are equipopular and $T(q)$ and $T(q_1)$ have the same spine structure.

$T(p_1)$ and $T(q_1)$ almost fit the conditions of Lemma 15: they are both left-justified and both have the same spine structure. However, we need to ensure that they both have the same length of topmost spine, and that length must be at least 2. The idea is to build a sort of scaffolding on top of each tree, use it to sort the rest of the spines, and then remove it at the end.

Let $p_2 = (p_1 \oplus 1) \ominus 1$, and consider the tree $T(p_2)$. This is a tree with the root node representing the second to last element in the in-order reading. The left subtree of the root is $T(p_1)$ and the right subtree is a single node with no children. The spine structure of $T(p_2)$ is the same as that of $T(p_1)$ with an additional spine of length 2 added at the top of the trunk. Similarly, let $q_2 = (q_1 \oplus 1) \ominus 1$. Then the spine structure of $T(q_2)$ is the same as that of $T(q_1)$ with an additional spine of length 2 on the top. Since $T(p_1)$ and $T(q_1)$ have the same spine structure, $T(p_2)$ and $T(q_2)$ have the same spine structure. So we invoke Lemma 15 on $T(p_2)$ and $T(q_2)$: both trees are left-justified, they have the same spine structure, and their topmost spines each have length 2. Thus p_2 and q_2 are equipopular.

Since $p_2 = (p_1 \oplus 1) \ominus 1$ and $q_2 = (q_1 \oplus 1) \ominus 1$, by Theorem 8, p_1 and q_1 are equipopular, and thus p and q are equipopular as desired. $\square$

We have computed the popularity for patterns of length $k \leq 7$ and a few n , and thus proved the following corollary.

Corollary 17. *If p and q are equipopular length- k patterns, with $k \leq 7$, then $T(p)$ and $T(q)$ have the same spine structure.*

Proof. A Python program computing the popularity of patterns p of length $k \leq 7$ confirms that, for all p and q with different spine structures, there is a sufficiently large n such that $A_n(p) \neq A_n(q)$. If n_k is the least integer such that $A_{n_k}(p) \neq A_{n_k}(q)$ for any two length- k patterns p and q with different spine structures, we have $n_4 = 5$, $n_5 = 6$, $n_6 = 8$, and $n_7 = 9$. $\square$

5 Conclusion and Future Work

We have shown that if the trees $T(p)$ and $T(q)$ have the same spine structure, then p and q are equipopular. The converse remains open. The results of this paper allow us to divide length- k patterns into a relatively small number of equivalence classes based on the spine structure of the tree they map to. However, the possibility remains that some of those equivalence classes may be equivalent to each other.

We have experimentally determined the popularity of all patterns of length $k \leq 7$ for relatively small n . We've computed enough values for n to determine that all of the popularity classes predicted by the main theorem are distinct. In other words, for patterns of length up to 7 the main theorem of this paper completely classifies 132-avoiding permutations into equipopularity classes. Based this evidence, we conjecture that the equivalence classes proved in this paper are the smallest possible in general.

Conjecture 18. If p and q are equipopular, then $T(p)$ and $T(q)$ have the same spine structure.

In other words, if p and q map to partitions of different shape, then they are not equipopular and thus the equivalence classes corresponding to partitions of different shapes are different.

Corollary 19. *If Conjecture 18 is true, then the number of equivalence classes for length- k patterns is the number of partitions of k .*

Proof. The conjecture states that each spine structure corresponds to exactly one equipopularity class. But a spine structure of a tree for a length- k pattern is just a partition of k . Since every possible partition of k can be a spine structure of a tree $T(p)$, the number of equivalence classes is the number of partitions of k . $\square$

We also observe that, in our experiments, the relative order of the popularity of equipopularity classes corresponds to a partial ordering of the spine structures of the trees, called the *refinement order*. Since a spine structure can be viewed as a partition, we state this order in terms of partitions.

A partition $\lambda_a = \{a_1, a_2, \dots, a_r\}$ is less than the partition $\lambda_b = \{b_1, b_2, \dots, b_s\}$ in refinement order if λ_b can be obtained from λ_a by combining parts of λ_a , that is, if each b_i is a sum of one or more a_j 's such that each a_j is used exactly once.

Example 20. The partition $\{2, 1, 1\}$ of 4 is less than $\{2, 2\}$ in refinement order, because the two 1's can be combined into a 2. $\{2, 1, 1\}$ is also less than $\{3, 1\}$ in refinement order, by combining a 1 and a 2 to obtain the 3. The partitions $\{2, 2\}$ and $\{3, 1\}$ are incomparable in this partial order.

We conjecture that the order of the popularity of the equivalence classes follows the refinement order of the spine structures of their trees.

Conjecture 21. Given patterns p and q , $A_n(p) \leq A_n(q)$ for all n if and only if the spine structure of $T(p)$ is less than or equal to the spine structure of $T(q)$ in refinement order.

This conjecture is supported by our calculations of the popularity of patterns with length up to 7. It makes intuitive sense as well. Bóna [1] proved previously that the increasing pattern of length k is the least popular and the decreasing pattern is the most popular, for all k . The structure of 132-avoiding permutations suggests that the closer a pattern is to being decreasing, the more popular it should be, and vice versa. The spines in the tree $T(p)$ represent decreasing subsequences of the pattern p , so in imprecise terms, longer spines give rise to a pattern that is “closer” to being completely decreasing, which suggests that the pattern should be more popular.

The original question of Cooper was more broad: what can be said about the popularity of patterns in $S_n(p)$ for other p than 132? First, note that the other non-monotone length-3 patterns, 231, 312, and 213 are the reverse, complement, and reverse-complement of 132 respectively. So our result for $S_n(132)$ can be transformed into a symmetric result

for any of the other patterns by taking the reverse or complement of the above results. Homberger [4] has a result on the popularity of length-3 patterns in $S_n(123)$, which symmetrically applies to $S_n(321)$. Results on the popularity of longer patterns in $S_n(132)$, or popularity in $S_n(p)$ for patterns p of length greater than 3 would be interesting.

Acknowledgments

This research was performed at the University of Minnesota Duluth under the supervision of Joe Gallian and supported by the National Science Foundation (grant number DMS-1062709) and the National Security Agency (grant number H98230-11-1-0224). Many thanks to Eric Riedl, David Rolnick, and Adam Hesterberg for discussing the problem with me and reading the paper. Thanks also to all of the visitors and program participants in the 2012 Duluth REU.

References

- [1] Miklos Bóna. The Absence of a Pattern and the Occurrences of Another. *Discrete Mathematics & Theoretical Computer Science*, 14(2), 2010.
- [2] Miklos Bóna. Surprising symmetries in Objects Counted by Catalan Numbers. *Electronic Journal of Combinatorics*, 19(1), 2012.
- [3] Joshua Cooper. Combinatorial Problems I Like. <http://www.math.sc.edu/~cooper/combprob.html>, 2012
- [4] Cheyne Homberger. Expected Patterns in Permutation Classes. [arXiv:1206.0320v2](https://arxiv.org/abs/1206.0320v2) 12 Jul 2012.
- [5] Miklos Bóna, Bruce Sagan, Vincent Vatter. Frequency sequences with no internal zeros. *Adv. Appl. Math.*, 28 (2002) 395-420.