

Uniquely-Wilf classes

Michael Albert

Jinge Li

Department of Computer Science, University of Otago

received 12th Apr. 2019, revised 25th Sep. 2019, accepted 15th Oct. 2019.

Two permutations in a class are Wilf-equivalent if, for every size, n , the number of permutations in the class of size n containing each of them is the same. Those infinite classes that have only one equivalence class in each size for this relation are characterised provided either that they avoid at least one permutation of size 3, or at least three permutations of size 4.

1 Introduction

In the classical study of permutation patterns, two permutations π and σ are said to be *Wilf-equivalent* if the class, $\text{Av}(\pi)$, of permutations avoiding π and the class, $\text{Av}(\sigma)$, of permutations avoiding σ have the same enumeration sequence. Notational definitions are given in the next section. Arguably, one of the main results that initiated this area of study was that any two permutations of size 3 are Wilf-equivalent. Due to underlying symmetries this relation could have had at most two equivalence classes on permutations of size 3. It has since been discovered that among the permutations of size 4 there are only three Wilf-equivalence classes while, taking only symmetry into account, there could have been up to seven.

In this paper we are concerned with a relativised version of Wilf-equivalence. Rather than considering all the permutations that avoid π we consider only those that belong to some given class \mathcal{C} – so the original notion just corresponds to taking \mathcal{C} to be the class of all permutations.

Definition 1. *Let \mathcal{C} be a permutation class. Two permutations $\pi, \sigma \in \mathcal{C}$ are Wilf-equivalent relative to \mathcal{C} if the two classes $\mathcal{C} \cap \text{Av}(\pi)$ and $\mathcal{C} \cap \text{Av}(\sigma)$ have the same enumeration sequence. In this case we write $\pi \equiv_{\mathcal{C}} \sigma$.*

If \mathcal{C} is clear from context we sometimes omit mentioning it.

Let $(c_n)_{n \in \mathbb{N}}$ be the enumeration sequence of a class \mathcal{C} . If $\pi \in \mathcal{C}$ has size k then the enumeration sequence $(p_n)_{n \in \mathbb{N}}$ of $\mathcal{C} \cap \text{Av}(\pi)$ satisfies $p_n = c_n$ for $n < k$ and $p_k = c_k - 1$. Therefore, if π and σ are Wilf-equivalent relative to \mathcal{C} they must have the same size.

Definition 2. *Let \mathcal{C} be a permutation class. The Wilf-sequence of \mathcal{C} is the sequence $(w_n)_{n \in \mathbb{N}}$ where w_n is the number of $\equiv_{\mathcal{C}}$ -equivalence classes among the permutations of \mathcal{C} of size n .*

Obviously $w_n \leq c_n$ in general but it has been observed in multiple contexts that it is frequently the case that $w_n = o(c_n)$. We refer to this phenomenon as a *Wilf-collapse*. Since it is known that the partially ordered sets corresponding to permutation classes do not have infinite automorphism groups (Albert et al. (2015)) this cannot occur solely due to global symmetries of the containment relation within a pattern class. Our particular concern in this paper is with the ultimate Wilf-collapse:

Definition 3. A permutation class \mathcal{C} is uniquely-Wilf if the equivalence relation $\equiv_{\mathcal{C}}$ has only one equivalence class on permutations of size n in \mathcal{C} for all n .

In Albert and Li (2019) it was shown that the class $\text{Av}(312, 213)$ is uniquely-Wilf. It is the aim of this paper to explore the landscape of uniquely-Wilf classes with the ultimate (but unrealised) goal of classifying them completely.

2 Preliminaries

In this section we introduce our notation, definitions and make some basic observations. Notations chosen are somewhat standard in the field of classical permutation patterns so we shall be quite brief. A primer on this notation can be found in Bevan (2015) and for introductions to and/or surveys of the field we suggest the books Bóna (2012) and Kitaev (2011) and a survey article Vatter (2015).

We denote by \mathcal{S} the class of all finite permutations. For a set of permutations X we denote the set of permutations in X of size k by X_k and similarly use $X_{\leq k}$ etc. The enumeration sequence $(x_n)_{n \in \mathbb{N}}$ of X is defined by $x_k = |X_k|$.

The set \mathcal{S} carries a partial order \leq called *classical pattern containment*. Specifically, $\pi \leq \sigma$ if σ , considered as a sequence in one line notation, contains a subsequence order-isomorphic to π . There are eight automorphisms of \mathcal{S} as a partially ordered set generated by the operations on sequences of: reverse, complement (subtract each term for a permutation of size n from $n + 1$) and inverse. All the questions we consider respect these automorphisms.

If $\pi \leq \sigma$ we say that σ *involves* π , while if $\pi \not\leq \sigma$ we say that σ *avoids* π . Further we define:

$$\begin{aligned} \text{Inv}(\pi) &= \{\sigma \in \mathcal{S} : \pi \leq \sigma\} \\ \text{Av}(\pi) &= \{\sigma \in \mathcal{S} : \pi \not\leq \sigma\}. \end{aligned}$$

When $\pi \leq \sigma$ and $|\sigma| = |\pi| + 1$ we say that σ *covers* π .

The $\text{Av}(\pi)$ notation is extended to sets of permutations with the requirement that all the permutations in the set must be avoided. A set of permutations is a *class* if it is downwards-closed with respect to \leq or equivalently if it is of the form $\text{Av}(X)$ for some set of permutations X . As already seen we generally use Greek letters to denote individual permutations, upper case Latin letters to denote sets of permutations, and calligraphic upper case Latin letters to denote permutation classes. Additionally we set $\mathcal{I} = \text{Av}(21)$ to be the class of all monotone-increasing permutations and $\mathcal{D} = \text{Av}(12)$ to be the class of all monotone-decreasing permutations.

Definition 4. Let k and n be positive integers. A permutation class \mathcal{C} is (k, n) -balanced if, for all $\pi, \sigma \in \mathcal{C}_k$, $|\mathcal{C}_n \cap \text{Inv}(\pi)| = |\mathcal{C}_n \cap \text{Inv}(\sigma)|$.

Any permutation class is (k, n) -balanced for any $k \geq n$ (since the two sets in question are singletons if $k = n$ and empty otherwise) so the issue of whether or not a class is (k, n) -balanced is really of interest only if $k < n$.

Observation 5. A permutation class \mathcal{C} is uniquely-Wilf if and only if it is (k, n) -balanced for all pairs of positive integers (k, n) .

Proof: Suppose that \mathcal{C} is uniquely-Wilf and that $\pi, \sigma \in \mathcal{C}_k$. Then $\pi \equiv_{\mathcal{C}} \sigma$ so the two classes $\mathcal{C} \cap \text{Av}(\pi)$ and $\mathcal{C} \cap \text{Av}(\sigma)$ have the same enumeration sequence. In particular $|\mathcal{C}_n \cap \text{Av}(\pi)| = |\mathcal{C}_n \cap \text{Av}(\sigma)|$. But

$\mathcal{C}_n \cap \text{Inv}(\pi) = \mathcal{C}_n \setminus (\mathcal{C}_n \cap \text{Av}(\pi))$ and similarly for σ so this is equivalent to $|\mathcal{C}_n \cap \text{Inv}(\pi)| = |\mathcal{C}_n \cap \text{Inv}(\sigma)|$, i.e., \mathcal{C} is (k, n) -balanced.

Conversely, suppose that \mathcal{C} is (k, n) -balanced for all pairs of positive integers (k, n) . Let $\pi, \sigma \in \mathcal{C}_k$ be given. Then, for any n ,

$$|\text{Av}(\pi) \cap \mathcal{C}_n| = c_n - |\text{Inv}(\pi) \cap \mathcal{C}_n| = c_n - |\text{Inv}(\sigma) \cap \mathcal{C}_n| = |\text{Av}(\sigma) \cap \mathcal{C}_n|,$$

i.e., $\pi \equiv_{\mathcal{C}} \sigma$, hence \mathcal{C} is uniquely-Wilf. \square

Observation 6. *Let \mathcal{C} be a uniquely-Wilf class that contains both 12 and 21. Then, for each k , either \mathcal{C} contains both $12 \cdots k$ and $k(k-1) \cdots 1$ or it contains neither of those permutations. In the latter case \mathcal{C} is finite.*

Proof: If, for some k , \mathcal{C} contained $12 \cdots k$ but not $k(k-1) \cdots 1$ then $12 \not\equiv_{\mathcal{C}} 21$ since \mathcal{C} contains a permutation of size k avoiding 21 but none avoiding 12. This contradiction establishes the first claimed result. The second part follows immediately since, by the Erdős-Szekeres theorem (Erdős and Szekeres (1935)), a class containing neither $12 \cdots k$ nor $k(k-1) \cdots 1$ cannot contain any permutation of size greater than $(k-1)^2$. \square

Observation 6 establishes a trichotomy for uniquely-Wilf classes: they might be finite, consist only of monotone permutations of a single type, or contain both monotone permutations of every size. Somewhat disingenuously we consider the first two cases as trivial ones and focus on the last. This is disingenuous since any uniquely-Wilf class \mathcal{C} is the union of all the finite uniquely-Wilf classes $\mathcal{C}_{\leq k}$ for each positive integer k . This observation is central to the experimental results described in the next section.

3 Experimental results

Let \mathcal{F} be a finite uniquely-Wilf class whose permutations are all of size at most n . We are interested in determining those uniquely-Wilf classes \mathcal{C} with the property that $\mathcal{C} \cap \mathcal{S}_{\leq n} = \mathcal{F}$. This condition already puts some restrictions on \mathcal{C} – specifically that any permutation of size at most n that does not belong to \mathcal{F} is avoided by all the elements of \mathcal{C} , i.e., $\mathcal{C} \subseteq \text{Av}(\mathcal{S}_{\leq n} \setminus \mathcal{F})$. For convenience we denote $\text{Av}(\mathcal{S}_{\leq n} \setminus \mathcal{F})$ by \mathcal{F}^\uparrow – this is the maximum class (in terms of set containment) whose restriction to $\mathcal{S}_{\leq n}$ is \mathcal{F} .

Observation 7. *Let \mathcal{C} be a permutation class and let k be a positive integer. Then $\mathcal{C} = (\mathcal{C}_{\leq k})^\uparrow$ if and only if $\mathcal{C} = \text{Av}(X)$ for some $X \subseteq \mathcal{S}_{\leq k}$.*

Proof: This is essentially a restatement of the definition. If $\mathcal{C} = (\mathcal{C}_{\leq k})^\uparrow$ then we can take X to be $\text{Av}(\mathcal{S}_{\leq k} \setminus \mathcal{C}_{\leq k})$. Conversely if $\mathcal{C} = \text{Av}(X)$ for some $X \subseteq \mathcal{S}_{\leq k}$ then $X \subseteq \mathcal{S}_{\leq k} \setminus \mathcal{C}_{\leq k}$ and so $\text{Av}(\mathcal{S}_{\leq k} \setminus \mathcal{C}_{\leq k}) \subseteq \mathcal{C}$. But the reverse containment always holds so $\mathcal{C} = \text{Av}(\mathcal{S}_{\leq k} \setminus \mathcal{C}_{\leq k})$. \square

If $\mathcal{F} \subseteq \mathcal{S}_{\leq n}$ is a finite uniquely-Wilf class, then an obvious question to address is:

Which subsets $X \subseteq \mathcal{S}_{n+1}$ have the property that $\mathcal{F} \cup X$ is a uniquely-Wilf class?

Call such subsets *potential extensions* of \mathcal{F} to size $n+1$. It's quite easy to determine whether X is a potential extension of \mathcal{F} of size $n+1$. First of all X must not contain any permutations that cannot lie in any extension of \mathcal{F} , i.e., $X \subseteq \mathcal{F}^\uparrow$. This ensures that $\mathcal{F} \cup X$ is a class. Secondly, $\mathcal{F} \cup X$ must

be $(k, n + 1)$ -balanced for all $k \leq n$. Since \mathcal{F} was already (k, m) balanced for all $m \leq n$ and $\mathcal{F} \cup X$ contains no permutations of size greater than $n + 1$ these two conditions are both necessary and sufficient for X to be a potential extension of \mathcal{F} .

This suggests an experimental technique to build uniquely-Wilf classes from the bottom up. Start with a finite permutation class \mathcal{F} that is uniquely-Wilf and contains permutations up to size n . Consider binary vectors (i.e., indicator functions) $\mathbf{x} : (\mathcal{F}^\dagger)_{n+1} \rightarrow \{0, 1\}$ where, as usual we denote $\mathbf{x}(\pi)$ by \mathbf{x}_π and associate \mathbf{x} with the set $X = \mathbf{x}^{-1}(1)$. Then *par abus de langage* we call \mathbf{x} a potential extension of \mathcal{F} to size $n + 1$ if the associated set X is. The important thing is that this construction already captures the first condition above since we only consider permutations from \mathcal{F}^\dagger . Each of the balancing conditions that any two permutations of the same size, say $\sigma, \tau \in \mathcal{F}_k$, have the same number of extensions in X can be expressed as a constraint:

$$\sum_{\sigma \leq \pi} \mathbf{x}_\pi - \sum_{\tau \leq \pi} \mathbf{x}_\pi = 0.$$

Here, and in the following few equations, π ranges over (\mathcal{F}^\dagger) (i.e., the set of permutations on which the notation \mathbf{x}_π is defined).

From a practical standpoint we might wish to change the form in which we express these constraints. For instance we could use instead:

$$\sum_{\sigma \leq \pi, \tau \not\leq \pi} \mathbf{x}_\pi - \sum_{\tau \leq \pi, \sigma \not\leq \pi} \mathbf{x}_\pi = 0.$$

Or, for each $2 \leq k \leq n$ we could assign a target, t_k , with

$$0 \leq t_k \leq \left| (\mathcal{F}^\dagger)_{n+1} \right|$$

and, for $\sigma \in \mathcal{F}_k$, use the constraint

$$\sum_{\sigma \leq \pi} \mathbf{x}_\pi = t_k.$$

Or we could simply look for all solutions that give $(n, n + 1)$ -balance (possibly with a target) and then filter out those that do not also give $(k, n + 1)$ -balance for $k < n$.

Whatever our chosen system of constraints might be it can then be passed to a constraint solver. We used the system Minion (Gent et al. (2006)). All solutions can be found and then (if necessary) these can be filtered to provide all the potential extensions. Then we can proceed in the same manner on each of these (trying to further extend to size $n + 2$) until either some resolution is achieved or we exhaust our patience or machine resources. What sorts of resolution can we hope for?

- We might find that none of the potential extensions of \mathcal{F} contain the monotone permutations of size $n + 1$. In this case we know that there is no infinite uniquely-Wilf class whose restrictions to permutations of size at most n is \mathcal{F} .
- We might find that the only potential extension of \mathcal{F} that contains both monotone permutations of size $n + 1$ is $\mathcal{F}_{n+1}^\dagger$. If this continues to be the case when we extend to $n + 2$, $n + 3$ etc. then we might form the opinion that \mathcal{F}^\dagger is uniquely-Wilf and is the only such class whose restriction to permutations of size at most n is \mathcal{F} . However, this does not constitute a proof of that opinion – but it is something we can attempt to prove.

Since we are interested only in infinite uniquely-Wilf classes containing both 12 and 21 our starting point is the finite class $\mathcal{S}_{\leq 2} = \{1, 12, 21\}$. The $(1, n)$ -balancing condition is always trivially satisfied and $(2, n)$ -balance is satisfied whenever we include both monotone permutations of size n . So our first sets of possible extensions are the union of $\{123, 321\}$ with each of the sixteen subsets of $\{132, 213, 231, 312\}$. Since all questions about Wilf-equivalence are invariant under the action of the automorphism group of (\mathcal{S}, \leq) we can further restrict our experiments to one representative of each orbit of that group on sets of permutations. Our findings are as follows:

Experimental results. *Let $\mathcal{F} \subseteq \mathcal{S}_{\leq 3}$ be a finite class with $\mathcal{F}_2 = \{12, 21\}$. Then:*

1. *If $|\mathcal{F}_3| = 2$ the only infinite uniquely-Wilf class containing \mathcal{F} is $\mathcal{F}^\uparrow = \mathcal{I} \cup \mathcal{D}$.*
2. *If $|\mathcal{F}_3| = 3$ the only infinite uniquely-Wilf class containing \mathcal{F} is \mathcal{F}^\uparrow .*
3. *If $|\mathcal{F}_3| = 4$ and $\mathcal{F}_3 = \{123, 213, 312, 321\}$ or one of its symmetries then the only infinite uniquely-Wilf class containing \mathcal{F} is \mathcal{F}^\uparrow while, for all other choices of \mathcal{F}_3 with $|\mathcal{F}_3| = 4$, there are no such classes.*
4. *If $|\mathcal{F}_3| = 5$ then there is no infinite uniquely-Wilf class containing \mathcal{F} .*
5. *If $\mathcal{F}_3 = \mathcal{S}_3$ and $|\mathcal{F}_4| < 22$ then there are no infinite uniquely-Wilf classes containing \mathcal{F} .*

In the next section we prove the first and second results above along with the first half of the third result. Provided that we trust our constraint solver (which we do) and our translation of the constraints into code⁽ⁱ⁾ for the constraint solver (which we have to) the second half of the third result, the fourth result, and the fifth result require no further proof.

Of note is that we have not completely classified infinite uniquely-Wilf classes since the cases where $\mathcal{F}_3 = \mathcal{S}_3$ and $|\mathcal{F}_4| \in \{22, 23, 24\}$ remain open – this is where exhaustion set in as there were simply too many possible extensions to deal with using our present ideas. While the case $|\mathcal{F}_4| = 22$ could probably be handled by our current techniques using some additional computational power it does seem that the remaining cases are definitely out of reach using the straight forward constraints presented here. In order to complete the classification would seem to require either determining some new constraints that have been overlooked, some clever ordering of the CSP that facilitates its (non-)solution, or a new structural idea.

However, we are confident enough to present:

Conjecture 8. *There are no infinite uniquely-Wilf classes, \mathcal{C} , with $\mathcal{C}_3 = \mathcal{S}_3$.*

4 Proofs

In this section we provide the required proofs that justify the resolutions of the experiment above.

Proposition 9. *The only infinite uniquely-Wilf class, \mathcal{C} with $\mathcal{C}_3 = \{123, 321\}$ is $\mathcal{I} \cup \mathcal{D}$.*


⁽ⁱ⁾ This code is freely available from the authors on request. Owing to its specialised nature we do not feel it appropriate, at this time, to make it part of a more widely available repository.

Proof: Certainly $\mathcal{I} \cup \mathcal{D}$ is uniquely-Wilf. If $\mathcal{C}_3 = \{123, 321\}$ then $\mathcal{C}_{\leq 3}^\uparrow = \mathcal{I} \cup \mathcal{D}$ as any non-monotone permutation of size 4 contains at least one non-monotone permutation of size 3. But any infinite uniquely-Wilf class containing 12 and 21 must contain $\mathcal{I} \cup \mathcal{D}$. \square

Up to symmetry, there is only one set of three permutations of size three containing $\{123, 321\}$ so the second result we need is:

Proposition 10. *The only infinite uniquely-Wilf class \mathcal{C} with $\mathcal{C}_3 = \{123, 132, 321\}$ is $\mathcal{C} = \text{Av}(213, 231, 312)$.*

Proof: Let $\mathcal{C} = \text{Av}(213, 231, 312)$. Since all the permutations avoided in \mathcal{C} are of size three, $\mathcal{C} = \mathcal{C}_{\leq 3}^\uparrow$. To verify that \mathcal{C} is uniquely-Wilf is easy. The permutations in \mathcal{C} have the structure shown in the picture below:




That is, they can be represented by an increasing sequence followed by a decreasing one. There is an obvious one to one correspondence between such permutations and pairs (a, b) of integers with $a \geq 0$ and $b \geq 1$. Here a represents the size of the increasing sequence and b the size of the decreasing one – and we always choose the maximum element to be in the decreasing sequence. Thus, the representation of the increasing permutation of size n in this form is as the pair $(n-1, 1)$. With respect to this representation $(a, b) \leq (c, d)$ as permutations if and only if $a \leq c$ and $b \leq d$ i.e., the containment order of permutations and the usual coordinatewise order coincide. In particular each permutation, (a, b) , of size n has exactly $k+1$ extensions of size $n+k$ namely those corresponding to pairs $(a+x, b+k-x)$ for $0 \leq x \leq k$. Thus, \mathcal{C} is uniquely-Wilf.

Note that, in \mathcal{C} every permutation has exactly two covers since (a, b) is covered by $(a+1, b)$ and by $(a, b+1)$.

Suppose that \mathcal{E} were a proper uniquely-Wilf subclass of \mathcal{C} containing all monotone permutations and with $\mathcal{E}_3 = \mathcal{C}_3$. Let $n+1$ be the least size of any permutation in $\mathcal{C} \setminus \mathcal{E}$ and note that $n \geq 3$. Since we omit one element of \mathcal{C} from \mathcal{E} in size $n+1$ there is an element of \mathcal{E} of size n with only one cover in \mathcal{E} . Therefore, as \mathcal{E} is uniquely-Wilf, every element of \mathcal{E} of size n must have exactly one cover of size $n+1$. In particular, the monotone permutation $(n, 1)$ must be the unique cover of $(n-1, 1)$ belonging to \mathcal{E} and hence its other cover $(n-1, 2) \notin \mathcal{E}$.

As $(n-2, 2) \in \mathcal{E}$, but $(n-1, 2) \notin \mathcal{E}$, it must be the case that $(n-2, 3) \in \mathcal{E}$ since that is the only remaining cover of $(n-2, 2)$ in \mathcal{C} . Now if $n=3$ we already have a contradiction since $(1, 3), (0, 4) \in \mathcal{E}$ are the two covers of $(0, 3) \in \mathcal{E}$. But if $n \geq 4$ then since $(n-3, 4)$ also covers $(n-3, 3)$ we have $(n-3, 4) \notin \mathcal{E}$. Since $(n-1, 2), (n-3, 4) \notin \mathcal{E}$ none of their covers can lie in \mathcal{E} . But these include $(n-1, 3)$ and $(n-2, 4)$ which are the two covers of $(n-2, 2)$. So $(n-2, 2)$ has no covers in \mathcal{E} and we have our contradiction (since \mathcal{E} is infinite all monotone permutations have one cover in \mathcal{E}). The latter part of the argument is perhaps better understood through the following picture where filled circles represent elements of \mathcal{E} and open ones elements of $\mathcal{C} \setminus \mathcal{E}$:




□

The second half of the proof above can also be adapted to characterise the finite uniquely-Wilf classes containing 123, 321 and one other permutation of size three. They agree with the maximum such class to some size, then include none or half (rounded up if necessary) of the permutations in that class at the next size, and then no further permutations.

Proposition 11. *Let $\mathcal{F} \subseteq \mathcal{S}_3$ be the finite class with $\mathcal{F}_3 = \{123, 132, 231, 321\}$. The only infinite uniquely-Wilf class containing \mathcal{F} is $\mathcal{F}^\uparrow = \text{Av}(213, 312)$.*

Proof:

That $\text{Av}(213, 312)$ is uniquely-Wilf is proven in Albert and Li (2019) so we provide only a short sketch of the argument here for the reader’s convenience. Permutations in this class have the wedge-shaped structure shown below:



The permutations of size n in the class are naturally encoded by binary sequences of length $n - 1$ over the alphabet $\{L, R\}$ obtained by reading the permutation from bottom to top and noting which “side” of the picture the corresponding element belongs. The containment relation in the class is exactly the subsequence containment relation for these words. Consider two words αL and βR where α and β are words of the same length. A word contains αL but none of its prefixes do if and only if it is of the form $AR^k L$ where A contains α but none of its prefixes do and $k \geq 0$. But similarly a word contains βR but none of its prefixes do if and only if it is of the form $BL^k R$ where B contains β but none of its prefixes do. This sets up a length-preserving bijection between the minimal words containing αL and those containing βR whenever we have such a bijection with respect to α and β . But of course such a bijection exists when α and β are both the empty word, so induction shows that this bijection exists in general for any two words of the same length. This bijection then extends to a length-preserving bijection between all words containing αL and all words containing βR since these are obtained from words of the first type by concatenation with an arbitrary word. This establishes that $\text{Av}(213, 312)$ is uniquely-Wilf.

Now suppose that \mathcal{E} is a proper subclass of $\text{Av}(213, 312)$ that contains \mathcal{F}_3 . The permutations in \mathcal{E} can be partitioned into their maximal monotone intervals which must be alternately increasing and decreasing when considered from bottom to top. There must be a bound on the number of such intervals (since \mathcal{E} is proper). Therefore the structure of a permutation in \mathcal{E} is determined (essentially) by a finite tuple of non-negative integers summing to its size and the enumeration of \mathcal{E} is polynomially bounded and hence ultimately polynomial (Kaiser and Klazar (2002/03)).

The structure of polynomial classes is well-understood and we will make use of a form introduced in Homberger and Vatter (2016) – we sketch only the necessary details here and refer to that paper for

the full results. Define a *peg permutation* to be a permutation written in one line notation where each of its elements is decorated with a superscript from the set $\{+, -, \bullet\}$. For example, $2^-3^-1^\bullet$ is a peg permutation.

Given a pegged permutation $\tilde{\rho}$ the set of permutations $\text{Grid}(\tilde{\rho})$ consists of all those permutation that can be obtained from ρ by inflating each point into a monotone (possibly empty) interval according to the associated symbol: for $+$ an arbitrary increasing interval, for $-$ an arbitrary decreasing interval, and for \bullet either a single point or nothing at all. For instance 4321, 2431 and 432651 all belong to $\text{Grid}(2^-3^-1^\bullet)$ but 123 does not.

Departing from the terminology of Homberger and Vatter (2016) we say that a pegged permutation is *properly pegged* if it contains no two element monotone interval of any of the following forms:

$$1^+2^+, 1^\bullet 2^+, 1^+2^\bullet, 2^-1^-, 2^\bullet 1^-, 2^-1^\bullet. \quad (1)$$

To be explicit, the first condition above means that there are no two consecutive elements a and $a + 1$ in the pegged permutation both of which are decorated with $+$. We define $\text{Grid}^f(\tilde{\rho})$ to be the subset of $\text{Grid}(\tilde{\rho})$ consisting of all inflations of ρ in which any symbol decorated with $+$ or $-$ is inflated into a monotone interval of size at least two, and any symbol decorated with \bullet is replaced by a single point (the superscript stands for “filled”). Restating the conclusion of Theorem 1.4 of Homberger and Vatter (2016) slightly:

There exists a finite set \tilde{H} of properly pegged permutations such that \mathcal{E} is the disjoint union of the sets $\text{Grid}^f(\tilde{\rho})$ for $\tilde{\rho} \in \tilde{H}$.

We remark that since \mathcal{E} is a class that, if $\tilde{\rho} \in \tilde{H}$ then in fact $\text{Grid}(\tilde{\rho}) \subseteq \mathcal{E}$ since $\text{Grid}(\tilde{\rho})$ is just the downward closure of $\text{Grid}^f(\tilde{\rho})$.

Choose $\tilde{\rho} \in \tilde{H}$ in such a way that the total number, m of $+$ and $-$ decoration on $\tilde{\rho}$ is maximal (among all the elements of \tilde{H}) and, among all such, the underlying permutation ρ is maximal in the pattern containment order. Equality is allowed in either case. For instance it might be the case that $2^+4^+1^\bullet 3^-$, $2^-4^\bullet 1^-3^+ \in \tilde{H}$. If no other permutation in \tilde{H} contains more than three non- \bullet decorations and no permutation with three non- \bullet decorations contains 2413 then either of these would be a suitable choice for $\tilde{\rho}$. Form a permutation $\pi \in \text{Grid}^f(\tilde{\rho})$ by inflating each such non- \bullet symbol to a monotone interval having at least K elements where K is greater than the square of the largest size of any permutation underlying an element of \tilde{H} . This choice of K ensures (easily) that the only covers of π in \mathcal{E} are obtained by extending the inflated intervals by one point. In particular π has exactly m covers in \mathcal{E} . Therefore, as \mathcal{E} is uniquely-Wilf, all sufficiently large permutations in \mathcal{E} must have exactly m covers.

Now suppose that $\text{Grid}(1^+3^+2^\bullet) \subseteq \mathcal{E}$. Then any permutation that consists of a monotone increasing sequence followed by one more element belongs to \mathcal{E} as these all occur as elements of $\text{Grid}(1^+3^+2^\bullet)$. In particular, for any n , the monotone increasing permutation of size n has at least $n + 1$ covers in \mathcal{E} which contradicts the observation above that all sufficiently large permutations in \mathcal{E} must have exactly m covers. So it follows by contradiction that $\text{Grid}(1^+3^+2^\bullet) \not\subseteq \mathcal{E}$. This, together with the fact that \mathcal{E} is a class, implies that no $\tilde{\rho} \in \tilde{H}$ contains a 132 pattern with both the 1 and the 3 carrying $+$ decorations – we’ll abbreviate this as “no 1^+3^+2 pattern”. Similarly, there can be no $\tilde{\rho} \in \tilde{H}$ containing a 23^-1^- pattern.

This technique of showing that certain restrictions on the labelings of the permutations in \tilde{H} must be satisfied because otherwise some permutation has “too many” covers contradicting the assumption that \mathcal{E} is uniquely-Wilf is used throughout the remainder of the argument.

The underlying permutations of \tilde{H} belong to $\text{Av}(312, 213)$ so if there were three or more non- \bullet decorations then two of them would occur on the same side and would therefore have the same sign. Since

the intervals described in (1) cannot occur they could not be consecutive, and this would create one of the forbidden situations above. It follows that there can be at most two non- \bullet decorations on each permutation of \tilde{H} and, if there are two, they must have differing signs.

Therefore, every sufficiently large permutation in \mathcal{E} has either exactly one or exactly two covers. The former case is easily dismissed – since the increasing permutation has only the increasing cover the only permutation in \tilde{H} carrying a $+$ decoration would have to be 1^+ and similarly 1^- would be the only one containing the $-$ decoration (and all the monotone permutations of \mathcal{E} would be represented by these along with 1^\bullet). The rest of the permutations (if any) would carry only \bullet s so define a finite set – but if any existed there would a non-monotone permutation with no covers. So $\mathcal{E} = \mathcal{I} \cup \mathcal{D}$ but supposedly $\mathcal{E}_3 = \mathcal{F}_3$ which contains non-monotone permutations.

Supposing that each sufficiently large permutation in \mathcal{E} has exactly two covers also implies that

$$\text{Grid}(2^+3^-1^\bullet) \not\subseteq \mathcal{E}$$

since otherwise the monotone increasing permutation of size n would then have all of:

$$123 \cdots (n+1), 23 \cdots (n+1)1, \text{ and } 123 \cdots (n-1)(n+1)n$$

as covers. Similarly, $\text{Grid}(1^\bullet3^+2^-) \not\subseteq \mathcal{E}$. Also,

$$\text{Grid}(1^+3^\bullet4^\bullet2^-) \not\subseteq \mathcal{E} \text{ and } \text{Grid}(2^+4^\bullet3^\bullet1^-) \not\subseteq \mathcal{E}$$

as otherwise the monotone increasing (or decreasing in the second case) permutation would have at least three extensions.

Putting all these observations together shows that the only possible elements of \tilde{H} containing both a $+$ and a $-$ decoration are 1^+2^- and 2^+1^- . Only one can occur as, if they both did, increasing permutations would have at least three covers in \mathcal{E} .

Suppose without loss of generality that $1^+2^- \in \tilde{H}$. We claim that, other than 1^+ , 1^- , and $1^\bullet2^-$ all other patterns in \tilde{H} contain only \bullet s as decorations.

To rule out other patterns containing symbols x^+ observe first that $\text{Grid}(2^+1^\bullet) \not\subseteq \mathcal{E}$ since if it were, there would be an extra extension to the monotone increasing permutation beyond the two already available from $\text{Grid}(1^+2^-)$. It follows, that the only possible $+$ decoration on an element of \tilde{H} can occur on the first (and least) element of a pattern. No decorations of 132 can belong to \tilde{H} since we can't place \bullet s on both the 3 and the 2 as that case is already covered by 1^+2^- and other labelings are forbidden by (1) or previous observations. By (1) we cannot begin 1^+2^\bullet . Finally, $\text{Grid}(1^+3^\bullet4^\bullet2^\bullet)$ also gives extra extensions to the monotone increasing permutation. Any larger labeled pattern includes one of these forbidden cases, and so we've ruled out the presence of a $+$ except for 1^+2^- and 1^+ .

Now we want to rule out patterns in \tilde{H} other than 1^- , $1^\bullet2^-$ and 1^+2^- that contain symbols x^- . We cannot appeal to symmetry because we've broken that by assuming $\text{Grid}(1^+2^-)$ is contained in \mathcal{E} . In $\text{Grid}(2^\bullet3^-1^\bullet)$ we gain an extra extension to the monotone decreasing permutation so no pattern in \tilde{H} can contain a 23^-1 pattern. Similarly no 231^- patterns can occur. Together these rule out all possibilities of having a symbol x^- other than the ones which are explicitly allowed.

Now $231 \in \mathcal{E}$ but is not in $\text{Grid}(1^+)$, $\text{Grid}(1^-)$, $\text{Grid}(1^\bullet2^-)$, or $\text{Grid}(1^+2^-)$ so its representation must be as $\text{Grid}(2^\bullet3^\bullet1^\bullet)$. The same applies for its supposed two covers in \mathcal{E} , each of their covers, and so on. However, \tilde{H} is finite, so we have our final contradiction. \square

The proof of Proposition 11 has obviously been the most intricate of the arguments so far. Fundamentally though it is not so different in character to the previous ones (indeed it would have been possible to absorb those into its proof). Working in a polynomial class is a powerful restriction because of the limit it imposes on the number of extensions of any sufficiently large permutation. This then enables the “surplus extensions” technique to rule out many configurations. The second author, as part of his current PhD research (to be submitted in 2020), has proven that the only uniquely-Wilf polynomial classes are the ones that already arise in Proposition 11.

5 Conclusion

Our goal of characterising the uniquely-Wilf permutation classes has not been fully realised as the possibility still exists that there may be some which we have not discovered. If so, they will contain all of \mathcal{S}_3 and at least 22 elements of \mathcal{S}_4 .

It is notable that most of the arguments by contradiction in Proposition 11 focused on the existence of surplus extensions to a monotone permutation. The monotone permutations are already somewhat special in the order of pattern containment since they are the only permutations that cover only one other permutation. Therefore it does not seem unreasonable to ask:

Question 12. *What can be determined about classes in which all non-monotone permutations are Wilf-equivalent?*

For polynomial classes, the second author has answered this question completely (forthcoming as part of his PhD thesis in 2020) and, other than the examples already discovered above there are only a few new examples for instance the class $\text{Av}(123, 312)$ which was considered in Albert and Li (2019).

The authors would like to acknowledge the helpful comments of two referees, which have improved the presentation of this work significantly.

References

- M. Albert and J. Li. Wilf-collapse in permutation classes having two basis elements of size three. *Austral. J. Comb.*, 73(3):440–455, 2019. ISSN 2202-3518.
- M. H. Albert, M. D. Atkinson, and A. Claesson. Isomorphisms between pattern classes. *J. Comb.*, 6(1-2): 3–17, 2015. ISSN 2156-3527. doi: 10.4310/JOC.2015.v6.n1.a1. URL <https://doi.org/10.4310/JOC.2015.v6.n1.a1>.
- D. Bevan. Permutation patterns: basic definitions and notation. *arXiv e-prints*, art. arXiv:1506.06673, Jun 2015.
- M. Bóna. *Combinatorics of permutations*. Discrete Mathematics and its Applications (Boca Raton). CRC Press, Boca Raton, FL, second edition, 2012. ISBN 978-1-4398-5051-0. doi: 10.1201/b12210. URL <https://doi.org/10.1201/b12210>. With a foreword by Richard Stanley.
- P. Erdős and G. Szekeres. A combinatorial problem in geometry. *Compos. Math.*, 2:463–470, 1935.
- I. P. Gent, C. Jefferson, and I. Miguel. Minion: A fast, scalable, constraint solver. In *Proceedings of the 2006 Conference on ECAI 2006: 17th European Conference on Artificial Intelligence August*

- 29 – September 1, 2006, Riva Del Garda, Italy, pages 98–102, Amsterdam, The Netherlands, The Netherlands, 2006. IOS Press. ISBN 1-58603-642-4. URL <http://dl.acm.org/citation.cfm?id=1567016.1567043>.
- C. Homberger and V. Vatter. On the effective and automatic enumeration of polynomial permutation classes. *J. Symbolic Comput.*, 76:84–96, 2016. ISSN 0747-7171. doi: 10.1016/j.jsc.2015.11.019. URL <https://doi.org/10.1016/j.jsc.2015.11.019>.
- T. Kaiser and M. Klazar. On growth rates of closed permutation classes. *Electron. J. Combin.*, 9(2): Research paper 10, 20, 2002/03. ISSN 1077-8926. URL http://www.combinatorics.org/Volume_9/Abstracts/v9i2r10.html. Permutation patterns (Otago, 2003).
- S. Kitaev. *Patterns in permutations and words*. Monographs in Theoretical Computer Science. An EATCS Series. Springer, Heidelberg, 2011. ISBN 978-3-642-17332-5; 978-3-642-17333-2. doi: 10.1007/978-3-642-17333-2. URL <https://doi.org/10.1007/978-3-642-17333-2>. With a foreword by Jeffrey B. Remmel.
- V. Vatter. Permutation classes. In *Handbook of enumerative combinatorics*, Discrete Math. Appl. (Boca Raton), pages 753–833. CRC Press, Boca Raton, FL, 2015.