CLASSIFICATION OF WILF-EQUIVALENCES FOR LENGTH 8 SINGLETON PATTERNS

ALEXANDER DOBNER

Given $\tau \in S_k$, a permutation $\sigma \in S_n$ *avoids* τ if there is no subsequence of σ whose elements are in the same relative order as τ. (Here we think of both τ and σ as ordered lists of numbers.) For example, if $\tau = 213$ then σ avoids τ if there is no $1 \le a < b < c \le n$ such that $\sigma(b) < \sigma(a) < \sigma(c)$. In this context τ is often referred to as a *pattern*. Pattern avoidance can be defined in terms of permutation matrices. Saying that σ avoids τ is equivalent to saying that the permutation matrix of σ does not contain the permutation matrix of τ as a submatrix.

Given a collection of patterns $\Pi = \{\tau_1, \ldots, \tau_m\}$, we define

$$S_n(\Pi) := \{\sigma \in S_n : \sigma \text{ avoids } \tau \text{ for all } \tau \in \Pi\}.$$

We call this the set of Π-*avoiding patterns of length* n. If Π is a singleton set $\{\tau\}$ then we just write $S_n(\tau)$ to denote this set.

A central goal in the study of pattern avoiding permutations is to understand the size of $|S_n(\Pi)|$ for a given Π and n. Certain patterns are well understood. For example, for $\tau = 231$, Knuth showed $|S_n(\tau)| = C_n$, the nth Catalan number. In fact, it turns out that $|S_n(\tau)| = C_n$ for *all* $\tau \in S_3$. This suggests the following question: for which sets of patterns Π_1 and Π_2 is it the case that $|S_n(\Pi_1)| = |S_n(\Pi_2)|$ for all $n \in \mathbb{N}$? If Π_1 and Π_2 have this property, then they are said to be *Wilf-equivalent*. This is written as $\Pi_1 \sim \Pi_2$. For singleton sets of patterns $\{\tau\}$ and $\{\tau'\}$ we write $\tau \sim \tau'$.

Finding a general classification of pattern sets up to Wilf-equivalence is a difficult problem. However, there are many known examples of Wilf-equivalence. The simplest of these are the "trivial" Wilf-equivalences that come from symmetry. These symmetries are best understood using permutation matrices. There is a natural action of D_4 (the symmetries of the square) on permutation matrices/patterns which preserves the notion of pattern avoidance. Consequently, for any pattern set Π we can apply these symmetries to get up to 7 other pattern sets that are Wilf-equivalent to Π.

Other Wilf-equivalences are more difficult to come by. Restricting to the case of singleton sets, the list of all known Wilf-equivalences are as follows:

(1) (Stankova 1994 [4]) $1342 \sim 2413$.
(2) (Stankova & West 2002 [3]) $231 \oplus \tau \sim 312 \oplus \tau$ for any permutation τ. (Here $\oplus$ refers to the *direct sum* of permutations. The quickest definition is as follows: $\rho \oplus \pi$ is the permutation you get from taking the permutation matrices of ρ and π and putting them together into a block diagonal permutation matrix.)
(3) (Backelin, West, & Xin 2007 [1]) $12 \cdots k \oplus \tau$ and $k \cdots 21 \oplus \tau$ for any $k \in \mathbb{N}$ and any permutation τ.

Note that (2) and (3) are infinite families of Wilf-equivalences, whereas (1) is a "sporadic" example which does not fall into either of the infinite families. Naturally we can ask whether these form a complete list.

Question 1. Are there other Wilf-equivalences for singleton pattern sets aside from those that are generated by the trivial equivalences and (1), (2), and (3)?

Date: November 12, 2025.

I am not sure whether anyone has conjectured an answer to this question, but it is natural to look for computational evidence one way or the other. To do this, we can try to search for all Wilf-equivalences among patterns of length k for small values of k.[1]

The problem of searching for all Wilf-equivalences is essentially the same thing as classifying singleton pattern sets up to Wilf-equivalence. For any fixed k, a straightforward algorithm for carrying out this classification is as follows.

(1) Generate the list of all $k!$ permutations in S_k. Apply the known Wilf-equivalences to reduce this to a list of candidate equivalence classes. Our goal is to determine whether these candidate equivalence classes are indeed distinct.
(2) Select a representative from each candidate equivalence class. For each representative τ, compute $|S_n(\tau)|$ for $n = 1, 2, \ldots, N$ up to some threshold N.
(3) If the sequences computed in Step 2 are all distinct, then the classification is complete. If there are some duplicate sequences, then compute an additional term for each of the relevant sequences (i.e. increment N and repeat Step 2, but only do the computation for the duplicated sequences). Repeat this process until there are are no longer any duplicate sequences.

If this procedure terminates then this gives a complete classification of patterns of length k. If it does not terminate, then there must be new Wilf-equivalence which is not covered by any of the known cases.

In practice, carrying out the algorithm above is quite computationally intensive. For $k \leq 7$, this was completed by Stankova and West [3, Fig. 9]. In their paper they do not say how long the computation took, but I imagine that for computers in the early 2000s it would be quite a long time. Going up to $k = 8$ was presumably computationally infeasible.

Nowadays computers are several orders of magnitude more powerful, so I was able to carry out the $k = 8$ calculation in less than a day on a laptop. I also had the advantage of being able to use some very efficient code for calculating $|S_n(\tau)|$ which was written by William Kuszmaul [2]. This search did not yield any new Wilf-equivalences. That is, all the sequences for the various candidate equivalence classes (of which there were 4755) ended up being distinct. This means we now have the following table (where the last column is new).

Pattern length	$k=1$	$k=2$	$k=3$	$k=4$	$k=5$	$k=6$	$k=7$	$k=8$
Wilf classes	1	1	1	3	16	91	595	4755

To get a sense of how long the steps took, here is a summary.

- Initially I computed $|S_n(\tau)|$ for all 4755 patterns for $n \leq 12$. The computation for these took about 1.6 seconds per pattern.
- After this there were 1055 patterns whose sequences were not unique. For each of these I ran the computation for $n = 13$. These took about 21 seconds per pattern.
- After the $n = 13$ computation there were 8 remaining patterns whose sequences were not unique. Computing $n = 14$ for each of these took about 5 minutes per pattern. In the end they were all unique.

I will post the resulting data along with this note. For $k = 9$ there are 42,681 candidate Wilf-equivalence classes. Completing the classification procedure for these is well within the realm of possibility for modern computers (especially given that the

[1]Note: it's not hard to see that if τ and τ' have have different lengths then we cannot have $\tau \sim \tau'$.

problem is highly parallelizable). For $k = 10$ one is probably better off looking for a different approach.

References

1. Jörgen Backelin, Julian West, and Guoce Xin, *Wilf-equivalence for singleton classes*, Adv. in Appl. Math. **38** (2007), no. 2, 133–148. MR 2290807
2. William Kuszmaul, `https://github.com/williamkuszmaul/patternavoidance`.
3. Zvezdelina Stankova and Julian West, *A new class of Wilf-equivalent permutations*, J. Algebraic Combin. **15** (2002), no. 3, 271–290. MR 1900628
4. Zvezdelina E. Stankova, *Forbidden subsequences*, Discrete Math. **132** (1994), no. 1-3, 291–316. MR 1297387