Young classes of permutations

Michael Albert

Department of Computer Science
University of Otago, Dunedin
New Zealand
`malbert@cs.otago.ac.nz`

Abstract

We characterise those classes of permutations closed under taking subsequences and reindexing that also have the property that for every tableau shape either every permutation of that shape or no permutation of that shape belongs to the class. The characterisation is in terms of the dominance order for partitions (and their conjugates) and shows that for any such class there is a constant k such that no permutation in the class can contain both an increasing and a decreasing sequence of length k.

1 Introduction

The most well known notion of the *shape* of a permutation arises from the Robinson-Schensted-Knuth correspondence, whereby the shape of $\pi \in \mathcal{S}_n$ is taken to be the shape of either of the Young tableaux produced from it by the RSK-algorithm. On the other hand, we might think of the shape of π as being described by its graph, i.e. the set of points $\{(i, \pi(i)) \mid 1 \le i \le n\}$.

In the first context, subshapes can be defined in terms of various natural orderings on the set of partitions. In the second context, there is a natural notion of *subpermutation* which we can introduce. The formal definition is spelled out in the next section but informally, $\sigma \in \mathcal{S}_k$ is a subpermutation of $\pi \in \mathcal{S}_n$ if, there is some k element subset of $\{(i, \pi(i)) \mid 1 \le i \le n\}$ whose relative ordering with respect to both axes is the same as that of $\{(j, \sigma(j)) \mid 1 \le j \le k\}$.

It seems natural to investigate the manner in which these two concepts interact. Unfortunately, the basic answer is: not very well. In particular Adin and Roichman [1] began such an investigation and reported indicative counterexamples to some tempting false conjectures, as well as more positive results relating to rectangular and hook shapes. More recently, at the Permutation Patterns 2010 conference at Dartmouth College, Panova [4], [2] and Tiefenbruck [8] reported on results in this area when attention is restricted to permutations avoiding certain specific patterns.

Since the two concepts seem somewhat orthogonal with respect to *specific* permutations, it remains to consider whether or not they interact well with respect to

collections of permutations. Indeed Schensted's theorem (and Greene's generalisation of it) establish a family of such relationships. For example: the collection of all permutations not containing a decreasing sequence of length 3 is the same as the collection of all permutations whose tableaux have at most two rows.

A set of permutations closed under taking subpermutations is called a *permutation class* (or simply *class*). Motivated by the result mentioned above arising from Schensted's theorem, we investigate the question:

> Which classes, $\mathcal{C}$, of permutations have the property that if π and π' are two permutations of the same shape then either both, or neither belong to $\mathcal{C}$?

Such classes are precisely those closed in the natural equivalence relation on permutations of "having the same shape". Classes of this type will be called *Young classes*.

Greene's generalisation of Schensted's theorem provides a natural candidate for a condition on the set of shapes that may occur in a Young class, and the main result of this paper is to show that this condition is both necessary (which is the new result) and sufficient (which is essentially a corollary of the generalisation). The most notable consequence of our characterisation of these classes is that for any proper Young class there is a constant k such that no permutation in the class contains both increasing and decreasing subsequences of length k.

2   Definitions

We refer the reader to either [5] or [7] for more detailed discussion of the RSK-correspondence and related matters. In this section we will collect only the basic definitions and results needed for the remainder of this paper.

A *partition*, λ, is a weakly decreasing sequence of positive integers, i.e. $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_k)$ where $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_k > 0$. We say that λ is a partition of n, and write $|\lambda| = n$ where $n = \sum_i \lambda_i$. The individual values λ_i are called the *parts* of λ. When writing a partition as a sequence, repeated values may be represented by exponentiation. A partition is often represented by its *Young diagram*, an array of boxes of the appropriate lengths (we use the English or "top down" convention). For instance:

$$(3, 3, 2, 1) = (3^2, 2, 1) = \ \boxed{}$$

It is sometimes convenient to pretend that a partition may be extended by (any number of) trailing 0's. The *conjugate partition*, λ^* of a partition λ is the sequence given by:

$$\lambda_i^* = |\{j \,:\, \lambda_j \geq i\}|\,.$$

The *sum* $\lambda + \mu$ of two partitions is defined in the obvious way (its i^{th} element is $\lambda_i + \mu_i$), and their *conjugate sum* is defined by:

$$\lambda +^* \mu = (\lambda^* + \mu^*)^*.$$

On Young diagrams these operations correspond to concatenating the two partitions horizontally or vertically, and then shoving blocks leftwards or upwards to eliminate gaps.

If λ and μ are partitions, we say that μ *dominates* λ and write $\lambda \trianglelefteq \mu$ if, for all $k \geq 1$:

$$\sum_{i=1}^{k} \lambda_i \leq \sum_{i=1}^{k} \mu_i.$$

If both $\lambda \trianglelefteq \mu$ and $\lambda^* \trianglelefteq \mu^*$ then we say that μ *doubly dominates* λ and write $\lambda \sqsubseteq \mu$. Since conjugation reverses the dominance order on partitions of n , if $|\lambda| = |\mu|$ and $\lambda \sqsubseteq \mu$ then $\lambda = \mu$.

The *shape* of permutation $\pi \in \mathcal{S}_n$ is the partition λ of n obtained when constructing its Young tableaux via the RSK-correspondence. In this case we write $\text{sh}(\pi) = \lambda$. We also say that n is the *length* of π since we frequently consider π and the sequence of its values interchangeably.

A permutation σ *is involved in* π if the terms of some subsequence of the values of π occur in the same relative order as the sequence of values of σ. In this case we write $\sigma \preceq \pi$, and also say that σ is a *subpermutation* of π. For example, $231 \preceq 421635$ because of the subsequence 463. A *permutation class* is a set of permutations closed downwards under $\preceq$. A permutation class can also be described as the set of permutations that do not involve (or *avoid*) some given set of permutations X. In that case we write $\mathcal{C} = \text{Av}(X)$.

A *Young class* is a class of permutations such that for each partition λ either all the permutations of shape λ or none of them belong to the class. The prototypical examples of Young classes are the collection of all permutations not including an increasing subsequence of some fixed length a. By Schensted's theorem [6], the shapes of all such permutations have every part smaller than a, and the converse also holds.

The use of Greek letters for both partitions and permutations is unfortunately standard. We will attempt to minimise the resulting confusion by reserving λ and μ for partitions.

3   Greene's Theorem and double domination

Greene's generalisation [3] (or see Theorem 3.5.3 of [5]) of Schensted's theorem provides a link between the shape of a permutation π and the involvement of certain patterns in π.

Theorem 1. *For any permutation π of tableau shape λ, and any k, the longest subpermutation of π that can be written as a union of k increasing subsequences has*

length $\lambda_1 + \lambda_2 + \cdots + \lambda_k$, and the longest subpermutation of π that can be written as a union of k decreasing subsequences has length $\lambda_1^* + \lambda_2^* + \cdots + \lambda_k^*$.

The increasing subsequences that make up the subpermutations whose existence is guaranteed by Greene's theorem can of course be taken to be disjoint, but they cannot in general be taken to have lengths $\lambda_1, \lambda_2, \ldots, \lambda_k$. A simple counterexample is the permutation 348951267 where $\lambda_1 = 5$, and $\lambda_2 = 3$. However, the unique increasing subsequence of length 5 is 34567, and the remaining elements, 8912, contain no three element increasing sequence. On the other hand, we can find two disjoint four element increasing sequences: 3489 and 1267.

Note that the permutations that avoid $(k+1)k \cdots 321$ are precisely those which can be written as a union of k increasing subsequences. So $\lambda_1 + \lambda_2 + \cdots + \lambda_k$ is the maximum length of an element of $\mathrm{Av}((k+1)k \cdots 321)$ which is involved in π. Of course a dual statement applies with reference to the second part of the theorem and the class $\mathrm{Av}(123 \cdots (k+1))$.

The following corollary to Greene's theorem is immediate:

Corollary 2. *Let σ and π be permutations. If $\sigma \preceq \pi$ then $\mathrm{sh}(\sigma) \sqsubseteq \mathrm{sh}(\pi)$.*

Proof. If $\sigma \preceq \pi$ then, for any k the longest subsequence of π which can be written as the union of k increasing (decreasing) subsequences is at least as long as the corresponding subsequence of σ, since that subsequence occurs in π. $\square$

A strong converse to the corollary is easily seen to be false. For example $(2,2) \sqsubseteq (3,1,1)$ but the permutation 54123 of shape $(3,1,1)$ has no subpermutation of shape $(2,2)$. However, even the weaker converse "if $\lambda \sqsubseteq \mu$ then for some permutations σ and π with $\mathrm{sh}(\sigma) = \lambda$ and $\mathrm{sh}(\pi) = \mu$, $\sigma \preceq \pi$" is also false. For, take $\lambda = (2,2,2)$ and $\mu = (4,1,1,1,1)$. If σ has shape λ, then to extend it to a permutation having a longest increasing subsequence of length 4 a pair of elements in increasing order must be added. However, to extend it to a permutation whose longest decreasing subsequence has length 5 a pair of elements in decreasing order must be added. To accomplish both therefore requires adding at least three elements, and thus rules out having shape μ. On the other hand, we will see below that if $\lambda \sqsubseteq \mu$ and $|\mu| = |\lambda| + 1$ then there are such permutations σ and π. This will be the central result that leads to a characterisation of Young classes.

For our subsequent arguments we will need to characterize the covering pairs for the relation $\sqsubseteq$ on partitions. The main point is to show that only the obvious pairs (those whose sizes differ by 1) are covering pairs.

Proposition 3. *Suppose that $\lambda \sqsubseteq \mu$, and $|\mu| - |\lambda| > 1$. Then there exists μ' with $\lambda \sqsubseteq \mu' \sqsubseteq \mu$ and $|\mu| - |\mu'| = 1$.*

Proof. We begin by eliminating some trivial or irrelevant cases.

- If $\lambda_1 = \mu_1$ or $\lambda_1^* = \mu_1^*$, then the result follows inductively by ignoring the first parts of λ and μ (or of their conjugates).
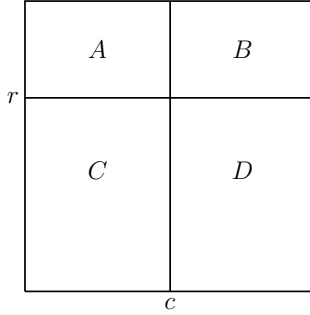
Figure 1: The subdivision of π and μ used in the proof of Proposition 3.

- If, for all r, $\sum_{i=1}^{r} \mu_i > \sum_{i=1}^{r} \lambda_i$, then we can form μ' by deleting the cell of the lowest row of the rightmost column of μ. This decreases each sum of row lengths by at most 1 (compared to μ) and leaves all but the final sum of column lengths unchanged. Since $\mu_1 > \lambda_1$ the only relevant comparisons of column sums occur before the final column, and so $\lambda \sqsubseteq \mu'$.

- A similar argument applies if, for all c, $\sum_{i=1}^{c} \mu_i^* > \sum_{i=1}^{c} \lambda_i^*$.

Assuming that none of the cases above hold, we are left with the following conditions: $\mu_1 > \lambda_1$; $\mu_1^* > \lambda_1^*$; for some $r > 1$, $\sum_{i=1}^{r} \mu_i = \sum_{i=1}^{r} \lambda_i$; and for some $c > 1$, $\sum_{i=1}^{c} \mu_i^* = \sum_{i=1}^{c} \lambda_i^*$. Choose the least possible values of r and c such that these equalities hold, and observe that $\mu_r < \lambda_r$ and $\mu_c^* < \lambda_c^*$.

Suppose that $\mu_r \geq c$ and $\mu_c^* \geq r$, i.e. that μ contains the upper left $r \times c$ rectangle. Since $\lambda_r > \mu_r$ and $\lambda_c^* > \mu_c^*$, λ also contains the upper left $r \times c$ rectangle. Divide the partitions λ and μ up into four regions as shown in Figure 1.

Both partitions contain the same number of cells in $A \cup B$ (since their r^{th} row sums agree) and also in $A \cup C$. Since both contain the entire region A, they contain the same number of cells in B and in C. So, there are at least two more cells of μ in region D than of λ, and the portion of μ in D doubly dominates the corresponding portion of λ. The result for this case now follows inductively by using it in D.

So we may now additionally assume that either $\mu_r < c$ or $\mu_c^* < r$ and hence that there is at least one square in the upper left $r \times c$ rectangle not belonging to μ. Form μ' by removing a cell from the rightmost column, and a cell from the bottommost row of μ and adding a cell within the upper left $r \times c$ rectangle. Then certainly $\mu' \sqsubseteq \mu$ and we claim that $\lambda \sqsubseteq \mu'$. Consider first the row sums of μ'. Some among the first $r - 1$ may be one smaller than the corresponding sums of μ but, by the choice of r they are still at least as large as those of λ. From rows r through to the final row they equal those of μ, and in the final row are again one smaller. However as $|\mu| > |\lambda|$ they are all at least as great as those of λ. The argument that the column sums of μ' dominate those of λ is exactly similar. $\square$

We now know that if $\lambda \sqsubseteq \mu$ is a covering pair then $|\mu| = |\lambda| + 1$. We can also bound the difference between row sums of λ and of μ. Specifically:

Proposition 4. *Suppose that $\lambda \sqsubseteq \mu$ and $|\mu| = |\lambda| + 1$. Then, for all r,*

$$0 \leq \sum_{i=1}^{r} \mu_i - \sum_{i=1}^{r} \lambda_i \leq 1.$$

Proof. The left hand inequality is part of the definition of $\lambda \sqsubseteq \mu$. For the remainder, suppose that it failed. Choose the least r such that $\sum_{i=1}^{r} \mu_i - \sum_{i=1}^{r} \lambda_i \geq 2$. Then $\mu_r > \lambda_r$. Let $c = \lambda_r$. Consider again the diagram in the preceding proof. Region A is fully occupied in both λ and μ, while region B contains at least two more cells of μ than of λ, and region D contains no cells of λ. So region C contains at least one more cell of λ than of μ but this (along with A being fully occupied) implies $\sum_{i=1}^{c} \lambda_i^* > \sum_{i=1}^{c} \mu_i^*$ and so contradicts $\lambda \sqsubseteq \mu$. $\square$

4   Young classes

To characterise Young classes we need to introduce two more operations on permutations, and show how they affect the shape of the resulting tableaux. Let permutations α and β of lengths n and k be given. Define:

$$\begin{aligned}
\alpha \oplus \beta &= (\alpha_1, \ldots, \alpha_n, n + \beta_1, \ldots, n + \beta_k) \\
\alpha \ominus \beta &= (k + \alpha_1, \ldots, k + \alpha_n, \beta_1, \ldots, \beta_k).
\end{aligned}$$

Informally, $\alpha \oplus \beta$ stacks β above and to the right of α, while $\alpha \ominus \beta$ stacks β below and to the right of α.

Proposition 5. *The operations $\oplus$ and $\ominus$ affect shape as follows:*

$$\begin{aligned}
\mathrm{sh}(\alpha \oplus \beta) &= \mathrm{sh}(\alpha) + \mathrm{sh}(\beta) \\
\mathrm{sh}(\alpha \ominus \beta) &= \mathrm{sh}(\alpha) +^* \mathrm{sh}(\beta).
\end{aligned}$$

Proof. The proof of both parts is most easily seen by considering Greene's theorem. In $\alpha \oplus \beta$ a subsequence that is the union of k increasing sequences can be divided into such a subsequence of α and such a subsequence of β (and the union of any two such subsequences is again of the same type). So, the sum of the first k parts of $\mathrm{sh}(\alpha \oplus \beta)$ is equal to the sum of the first k parts of $\mathrm{sh}(\alpha)$ plus the sum of the first k parts of $\mathrm{sh}(\beta)$. The result follows immediately. The corresponding result for $\ominus$ likewise follows by considering unions of decreasing subsequences.

$\square$

We remark that the result for $\oplus$ is also easily seen by considering the actual bumping operations of the RSK algorithm, but the result for $\ominus$ is not so transparent in that context (though it is true that the operations of RSK "push straight down columns").

Proposition 6. *Let* $\lambda \sqsubseteq \mu$ *be a covering pair for the double domination ordering.* *Then there exist permutations* σ *and* π *such that:* $\mathrm{sh}(\sigma) = \lambda$, $\mathrm{sh}(\pi) = \mu$ *and* $\sigma \preceq \pi$.

Proof. The proof is by induction on $|\lambda|$. The base case $|\lambda| = 0$ is completely trivial. Consider now a covering pair $\lambda \sqsubseteq \mu$.

Suppose first that for some i, $\lambda_i = \mu_i = k$, and choose i to be the least such. Let λ' and μ' be obtained by deleting the i^{th} part from λ and μ respectively. Inductively, choose $\sigma' \preceq \pi'$ with $\mathrm{sh}(\sigma') = \lambda'$ and $\mathrm{sh}(\pi') = \mathrm{sh}(\sigma')$. Let $\iota = 123\cdots k$, and let $\sigma = \iota \ominus \sigma'$ and $\pi = \iota \ominus \pi'$. Then $\sigma \preceq \pi$, and

$$\mathrm{sh}(\sigma) = (k) +^* \lambda' = \lambda$$
$$\mathrm{sh}(\pi) = (k) +^* \mu' = \mu.$$

Thus the result is true in this case, and by a parallel argument in the case where $\lambda_i^* = \mu_i^*$ for some i.

So we now assume that no part of λ is the same size as the corresponding part of μ (and that this also holds for their conjugates). Since $\lambda \sqsubseteq \mu$ and since we know that the row sums of μ can exceed the corresponding row sums of λ by at most 1, the row sums of μ must alternately exceed by 1 and equal the corresponding row sums of λ (otherwise we would have two equal parts). So:

$$\mu_1 = \lambda_1 + 1$$
$$\mu_2 = \lambda_2 - 1$$
$$\mu_3 = \lambda_3 + 1$$
$$\cdots$$

The partition μ must have exactly one more part than λ and this final part must be 1. All together this implies that for some positive integer k:

$$\mu = (2k+1, 2k-1, 2k-1, 2k-3, 2k-3, \ldots, 3, 3, 1)$$
$$\lambda = (2k, 2k, 2k-2, 2k-2, \ldots, 2, 2).$$

Since we shall need to do an induction on k, denote the first of these by $\mu^{(k)}$ and the second by $\lambda^{(k)}$. It remains to prove the result in precisely this case. The general construction is based on the first case of this sequence $(2,2) \sqsubseteq (3,1)$. Here we may take $\sigma = 2413$ and $\pi = 25314$. For use in the general case, define:

$$\theta^{(k)} = 1 \oplus 21 \oplus 321 \oplus \cdots \oplus k(k-1)\ldots 321.$$

The key property of $\theta^{(k)}$ is that its maximal increasing and decreasing subsequences both have length k and the deletion of either a maximal increasing or a maximal decreasing subsequence leaves a copy of $\theta^{(k-1)}$.

Let $\sigma^{(k)}$ be formed from $\sigma = 2413$ by "inflating" each of the elements into a copy of θ^k, and correspondingly let $\pi^{(k)}$ be obtained from $\pi = 25314$ by inflating the elements 2, 5, 1, and 4 in the same way, while leaving the point corresponding to 3
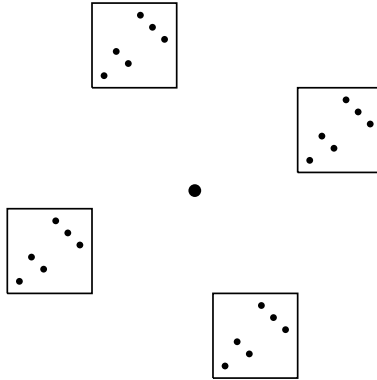
Figure 2: The permutation $\pi^{(3)}$ illustrating the proof of Proposition 6.

as a singleton. For instance, $\pi^{(3)} \in \mathcal{S}_{25}$ is shown in Figure 2 (and $\sigma^{(3)}$ is obtained from it by deleting the central point).

We claim that $\lambda^{(k)} = \mathrm{sh}(\sigma^{(k)})$ and $\mu^{(k)} = \mathrm{sh}(\pi^{(k)})$ (and this of course will finish the proof). This is proved by induction on k, with the case $k = 1$ already having been considered above.

We will provide the argument for $\pi^{(k)}$ in detail. Clearly an increasing sequence in $\pi^{(k)}$ can contain only elements in two of the inflated regions, together with possibly the central point (if the regions used are the leftmost and rightmost). In particular, the maximum increasing sequence has length $2k + 1$. However, if two increasing sequences include one using the central point then they can use only three of the four inflated regions, giving a maximum possible length of $4k$ (in two of the three regions we can only use an increasing sequence so obtaining $2k$ elements, in one we might use two increasing sequences giving a further $2k - 1$ elements, and finally we could include the central point). On the other hand we could also achieve a length of $4k$ by choosing an increasing sequence of length $2k$ from the two leftmost regions, and another one from the two rightmost regions. So, the first two rows of $\mathrm{sh}(\pi^{(k)})$ have lengths $2k + 1$ and $2k - 1$. In our second construction we found two increasing sequences in $\pi^{(k)}$ of total length $4k$ whose removal leaves a copy of $\pi^{(k-1)}$. By induction, the row sums of $\mathrm{sh}(\pi^{(k)})$ are at least as great as those of $\mu^{(k)}$. The same argument applied to decreasing sequences gives the corresponding result for column sums. So $\mu^{(k)} \sqsubseteq \mathrm{sh}(\pi^{(k)})$. Therefore, as the two partitions have the same size, they must be equal.

The argument for $\sigma^{(k)}$ is entirely similar, and in fact simplified by the omission of the central point. $\qquad\square$

Finally, we can prove our main result:

Theorem 7. *A permutation class $\mathcal{C}$ is a Young class if and only if there is some downwards closed set D in the double domination order of partitions and $\mathcal{C} = \{\pi : \mathrm{sh}(\pi) \in D\}$.*

Proof. By definition, $\mathcal{C}$ must be a union of shape equivalence classes, over a set D of shapes. To see that it is necessary that D be a downwards closed set in the double domination order suppose that $\mu \in D$. By Proposition 6, if $\lambda \sqsubseteq \mu$ has $|\lambda| + 1 = |\mu|$, there are two permutations σ and π with $\mathrm{sh}(\sigma) = \lambda$, $\mathrm{sh}(\pi) = \mu$ and $\sigma \preceq \pi$. As $\pi \in \mathcal{C}$ we must have $\sigma \in \mathcal{C}$ and hence $\lambda \in D$. However, by Proposition 3, the full double domination relation is the transitive closure of these single element covers. So, D must be downwards closed.

On the other hand, Corollary 2 says precisely that if D is downwards closed, then so is $\mathcal{C}$. $\qquad\square$

Note that for any partition λ there exists a hook, i.e. a partition of the form $(n, 1^m)$, such that $\lambda \sqsubseteq (n, 1^m)$ (essentially we can replace each element not in the first row or column by a pair of elements, one in the first row and one in the first column). It follows that in any proper Young class there are bounds, say a and d such that every permutation in the class either has no increasing subsequence of length a or no decreasing subsequence of length d. So every proper Young class is contained in the union of two of the prototypical Young classes already described by Schensted's theorem. Since, at the cost of enlarging one of those two classes, we could always take the monotone permutations which they avoid to have equal length, the final remark of the introduction is justified.

5   Conclusions and Open Problems

We have considered classes of permutations closed under the equivalence relation of having tableaux of the same shape. The general relationship between the involvement relationship for permutations and their shapes seems to be somewhat obscure. However, by restricting attention to these Young classes we were able to demonstrate a close connection with the double domination order for partitions.

More detailed connections (e.g. to describe the bases of these classes) seem to be difficult to achieve except in cases closely related to the prototypical Young classes provided by Schensted's theorem.

Returning to the vexing case of the relationship between involvement and shape for pairs of permutations (or tableau shapes), the most obvious remaining open question is:

> Give necessary and sufficient conditions on partitions $\lambda \sqsubseteq \mu$ which guarantee the existence of permutations σ and π with $\mathrm{sh}(\sigma) = \lambda$, $\mathrm{sh}(\pi) = \mu$ and $\sigma \preceq \pi$.

The example following Corollary 2 would seem to suggest that such conditions would involve a strengthening or refinement of Greene's theorem, but the details remain elusive.

Acknowledgments

References

[1] Ron M. Adin and Yuval Roichman, Shape avoiding permutations, *J. Combin. Theory Ser. A* **97** (2002), no. 1, 162–176.

[2] AndrewCrites, Greta Panova and Gregory S. Warrington, Separable permutations and Greene's theorem, (2010), available at `http://arxiv.org/abs/1011.5491`.

[3] Curtis Greene, An extension of Schensted's theorem, *Advances in Math.* **14** (1974), 254–265.

[4] Greta Panova, Separable permutations, Robinson-Schensted and shortest containing supersequences, *Permutation Patterns* (Dartmouth College, 2010), available at `http://www.math.dartmouth.edu/ pp2010/abstracts/Panova.pdf`.

[5] Bruce E. Sagan, *The symmetric group*, 2nd ed., Graduate Texts in Mathematics, vol. 203, Springer-Verlag, New York, 2001.

[6] C. Schensted, Longest increasing and decreasing subsequences, *Canad. J. Math.* **13** (1961), 179–191.

[7] Richard P. Stanley, *Enumerative combinatorics, Vol. 2*, Cambridge Studies in Advanced Mathematics, vol. 62, Cambridge University Press, Cambridge, 1999.

[8] Mark Tiefenbruck, 231-avoiding permutations and the Schensted correspondence, *Permutation Patterns* Dartmouth College (2010), available at `http://www.math.dartmouth.edu/ pp2010/abstracts/Tiefenbruck.pdf`.