

On intervals of the consecutive pattern poset

Sergi Elizalde^{1†} and Peter R. W. McNamara^{2‡}

¹*Department of Mathematics, Dartmouth College, Hanover, NH 03755, USA*

²*Department of Mathematics, Bucknell University, Lewisburg, PA 17837, USA*

Abstract. The consecutive pattern poset is the infinite partially ordered set of all permutations where $\sigma \leq \tau$ if τ has a subsequence of adjacent entries in the same relative order as the entries of σ . We study the structure of the intervals in this poset from topological, poset-theoretic, and enumerative perspectives. In particular, we prove that all intervals are rank-unimodal and strongly Sperner, and we characterize disconnected and shellable intervals. We also show that most intervals are not shellable and have Möbius function equal to zero.

Résumé. Le poset des motifs consécutifs est l'ensemble infini partiellement ordonné de toutes les permutations où $\sigma \leq \tau$ si τ a une sous-séquence des entrées adjacentes dans le même ordre relatif que les entrées de σ . Nous étudions la structure des intervalles dans ce poset des perspectives topologique et énumérative. En particulier, nous prouvons que tous les intervalles sont de rang unimodale et fortement Sperner, et nous caractérisons les intervalles déconnectés et enveloppables ("shellable"). Nous montrons également que la plupart des intervalles ne sont pas enveloppables et ils ont fonction de Möbius égal à zéro.

Keywords. Consecutive pattern, poset, disconnected, shellable, rank unimodal, Sperner, exterior, Möbius function

1 Introduction

Consecutive patterns in permutations generalize well-studied notions such as descents, ascents, peaks, valleys and runs. A permutation σ is said to be *contained* in another one τ as a *consecutive pattern* if τ has a subsequence of adjacent entries in the same relative order as the entries of σ . Otherwise, τ is said to *avoid* σ . For example, permutations that avoid both 123 and 321 as a consecutive pattern are the well-known alternating and reverse-alternating permutations. A systematic study of enumerative aspects of consecutive patterns started in [EN03] and, in the last decade, such patterns have become a vibrant area of research; see [Eli15] for a recent survey. Underlying all these questions is a partial order \mathcal{P} on the set of all permutations, where we define $\sigma \leq \tau$ if σ is contained in τ as a consecutive pattern. This paper is the first systematic study of intervals in this consecutive pattern poset. See Figure 1.1(a) for an example of such an interval.

Some questions for consecutive patterns are motivated by the analogous problems for so-called classical patterns, one of the most actively studied topics of combinatorics in the last three decades. For the

[†]Email: sergi.elizalde@dartmouth.edu. Partially supported by grant #280575 from the Simons Foundation, by grant H98230-14-1-0125 from the NSA, and by the Distinguished Visiting Professor Program at Bucknell University.

[‡]Email: peter.mcnamara@bucknell.edu. Partially supported by grant #245597 from the Simons Foundation.

definition of containment in this classical case, we remove the restriction that the relevant subsequence of τ consists of adjacent entries. For example, 123 is less than 2314 in the classical pattern poset but not in the consecutive pattern poset. See [Bón12] for an exposition of some of the most important developments in the study of classical patterns, [Kit11] for a detailed treatment, and [Ste13] for a survey of recent developments.

Examples of questions for consecutive patterns that have been motivated by the classical case include classifying them into equivalence classes determined by the number of permutations avoiding them, finding generating functions for the distribution of occurrences of a fixed pattern in permutations, and determining the asymptotic growth of the number of permutations avoiding a pattern. Consecutive patterns are interesting not only because the answers to the above questions and the techniques used to solve them are usually quite different than for classical patterns, but also because they have important applications to dynamical systems [AEK08].

While most of the aforementioned work is enumerative, our approach also has a poset-theoretic and topological flavor. An early inspiration for research from this viewpoint was a question of Wilf [Wil02] asking for the Möbius function of intervals in the poset defined by classical pattern containment. This question remains wide open, but has received increasing attention of late [BJJS11, MS15, SV06, Smi14, Smi15, ST10]. In contrast, the Möbius function for intervals in \mathcal{P} , the consecutive pattern poset, has been determined by Bernini, Ferrari and Steingrímsson [BFS11] and by Sagan and Willenbring [SW12]. This already gives an indication that the consecutive pattern case is more tractable for certain types of questions than the classical case.

The precursor in the classical case to the present work is [MS15], where the focus is on classifying disconnected open intervals and showing that certain special intervals are shellable. We successfully address these same topological questions of disconnectivity and shellability for the consecutive pattern poset \mathcal{P} . Furthermore, we consider what poset-theoretic properties might hold for \mathcal{P} , e.g., we show that all intervals are rank-unimodal, a statement that is just a conjecture in the classical case [MS15]. Our main results include the following.

- A simple classification of those open intervals in \mathcal{P} that are disconnected (Theorem 3.2).
- Almost all intervals contain a disconnected subinterval of rank at least 3, in a certain precise sense, and are thus not shellable (Theorem 3.6 and Corollary 4.1).
- All other intervals are shellable (Theorem 4.2). One motivation for shellability is that it completely determines the homotopy type of the order complex of the interval as a wedge of spheres. The number of spheres is the absolute value of the Möbius function and is thus readily determined using [BFS11, SW12].
- All intervals are rank-unimodal (Corollary 5.2).
- All intervals are strongly Sperner (Theorem 5.5), a condition equating the sizes of the union of the k largest ranks and the largest union of k antichains, for all k .
- It is clear from the formula for the Möbius function $\mu(\sigma, \tau)$ that it depends heavily on the *exterior* of τ , which is the longest permutation that is both a proper prefix and proper suffix of τ . We initiate a study of the asymptotic behavior of the exterior in Section 6. We show that its expected length is bounded between $e - 1$ and e (Theorem 6.4).

- Almost all intervals have a zero Möbius function, in a certain precise sense (Corollary 6.9).

Taken together, these topological, poset-theoretic and enumerative results give a rather comprehensive picture of the structure of intervals in \mathcal{P} .

Preliminaries and some initial observations are the content of the next section. Disconnectivity is the subject of Section 3, shellability is addressed in Section 4, and the properties of rank-unimodality and strongly Sperner are the topic of Section 5. The exterior of a permutation is considered in Section 6, along with several open problems about its behavior.

The proofs that have been omitted in this extended abstract due to lack of space can be found in the full paper version [EM15].




Fig. 1.1: The interval $[12, 213546]$ in \mathcal{P} (left) and its associated order complex (right).

2 Preliminaries

In this section, we collect together some useful terminology and notation, and make some initial observations.

2.1 Fundamentals

Let \mathcal{S}_n denote the set of permutations of $[n] := \{1, 2, \dots, n\}$. If $\tau \in \mathcal{S}_n$, we denote the length of τ by $|\tau| = n$, and we write τ in one-line notation as $\tau = \tau_1\tau_2 \dots \tau_n$. Given a sequence $a_1a_2 \dots a_k$ of distinct positive integers, define its *reduction* $\rho(a_1a_2 \dots a_k)$ to be the permutation of $\{1, \dots, k\}$ obtained by replacing the smallest entry with 1, the second smallest with 2, and so on. For example, $\rho(394176) = 263154$. For $1 \leq i \leq j \leq n$, let $\tau_{[i,j]} = \rho(\tau_i \dots \tau_j)$, and note that $\tau_{[i,j]} \in \mathcal{S}_{j-i+1}$.

We can now define our poset of interest. We say that τ *contains* σ as a *consecutive pattern* if there exist $1 \leq i \leq j \leq n$ such that $\tau_{[i,j]} = \sigma$. We write $\sigma \leq \tau$, and we say that $\tau_i \dots \tau_j$ is an *occurrence* of σ . The relation \leq defines a partial order on the set $\mathcal{S} = \bigcup_{n \geq 1} \mathcal{S}_n$ of all permutations, and we denote by \mathcal{P} the corresponding partially ordered set, called the *consecutive pattern poset*. The poset \mathcal{P} is the main object of study in this paper. See Figure 1.1(a) for an example of an interval in \mathcal{P} .

A related partial order on \mathcal{S} is obtained by considering classical pattern containment instead, i.e., the case when τ is said to contain σ if any subsequence of τ (not necessarily in consecutive positions) has reduction equal to σ . We will refer to the resulting poset as the *classical pattern poset*.

In the rest of this paper we will say that τ *contains* σ to mean that τ contains σ as a consecutive pattern, that is, $\sigma \leq \tau$. If this is not the case, we say that τ *avoids* σ .

To consider disconnectivity in Section 3 and shellability in Section 4, we need to define a simplicial complex associated with any interval of any poset. Given an interval $[\sigma, \tau]$, its *order complex* $\Delta(\sigma, \tau)$ is the simplicial complex whose faces are the chains of the open interval (σ, τ) . For example, the order complex of the interval $[12, 213546]$ is shown in Figure 1.1(b).

2.2 The Möbius function papers

It remains a wide open problem to determine the Möbius function for the *classical* pattern poset; see [BJJS11, MS15, SV06, Smi14, Smi15, ST10] for special cases. In contrast, the Möbius function for the consecutive pattern poset \mathcal{P} has been determined independently by Bernini, Ferrari and Steingrímsson in [BFS11] and by Sagan and Willenbring in [SW12]. We will present this recursive formula as stated in the latter paper here because it introduces some key ideas we will need later, and then we will explain the parts of the viewpoint of the former paper that will be useful to us.

We say that σ is a *prefix* (respectively *suffix*) of $\tau \in \mathcal{S}_n$ if $\sigma = \tau_{[1,i]}$ (resp. $\sigma = \tau_{[n-i+1,n]}$) for some i , and is a *proper prefix* (resp. *suffix*) if it does not equal τ . We say σ is a *bifix* of τ if it is both a proper prefix and proper suffix of τ . The *exterior* of τ , denoted $x(\tau)$, is the longest bifix of τ . Note that $x(\tau)$ always exists whenever $\tau > 1$ since 1 is a proper prefix and suffix of τ in that case. The *interior* of τ is $\tau_{[2,n-1]}$ and is denoted $i(\tau)$. For example, if $\tau = 21435$, then $x(\tau) = 213$ and $i(\tau) = 132$. We say that τ is *monotone* if it equals $12 \dots n$ or $n \dots 21$.

Theorem 2.1 ([BFS11, SW12]) For $\sigma, \tau \in \mathcal{P}$ with $\sigma \leq \tau$,

$$\mu(\sigma, \tau) = \begin{cases} \mu(\sigma, x(\tau)) & \text{if } |\tau| - |\sigma| > 2 \text{ and } \sigma \leq x(\tau) \not\leq i(\tau), \\ 1 & \text{if } |\tau| - |\sigma| = 2, \tau \text{ is not monotone, and } \sigma \in \{i(\tau), x(\tau)\}, \\ (-1)^{|\tau| - |\sigma|} & \text{if } |\tau| - |\sigma| < 2, \\ 0 & \text{otherwise.} \end{cases}$$

It is clear from this important theorem that the exterior of a permutation is worthy of attention, and yet this attention seems to be absent from the literature. This motivates our detailed consideration of the statistic $|x(\tau)|$ for general permutations in Subsections 6.1 and 6.2. In particular, we show that the expected value of $|x(\tau)|$ as $|\tau| \rightarrow \infty$ is bounded between $e - 1$ and e .

In [BFS11], if the exterior $x(\tau)$ satisfies the condition that $\sigma \leq x(\tau) \not\leq i(\tau)$, then $x(\tau)$ is called the *carrier element* of $[\sigma, \tau]$. As we see in Theorem 2.1, when $[\sigma, \tau]$ has a carrier element we get the interesting case where the calculation of $\mu(\sigma, \tau)$ reduces to determining $\mu(\sigma, x(\tau))$. It is therefore natural to ask under what conditions $[\sigma, \tau]$ has a carrier element. This motivates our consideration in Subsection 6.3 of the permutations τ such that $[1, \tau]$ has a carrier element. In this case of $\sigma = 1$, $x(\tau)$ can be called the *carrier element* of τ , and it exists if and only if $x(\tau) \not\leq i(\tau)$. We determine the asymptotic probability of a permutation not having a carrier element and, as an application, prove that the Möbius function is almost always 0, in a particular precise way.

2.3 Initial observations

One feature of the consecutive pattern poset that makes it more tractable than the classical pattern poset is that any permutation τ of length n covers at most two elements, namely $\tau_{[1,n-1]}$ and $\tau_{[2,n]}$. The next lemma will be useful later; it appears in [BFS11, SW12] and is routine to check.

Lemma 2.2 *For $\tau \in \mathcal{P}$ of length n , we have $\tau_{[1,n-1]} = \tau_{[2,n]}$ if and only if τ is monotone.*

As a consequence, we get our first structural result about intervals in \mathcal{P} .

Proposition 2.3 *The interval $[\sigma, \tau]$ in \mathcal{P} is a chain if and only if either*

- τ is monotone, or
- σ occurs exactly once in τ and does so as a prefix or a suffix.

As previously shown in [BFS11], we can completely determine the structure of $[\sigma, \tau]$ when σ has just one occurrence in τ , a result we state in Proposition 2.4. This contrasts with the classical pattern poset, where the corresponding result is unknown.

Proposition 2.4 ([BFS11]) *For an interval $[\sigma, \tau]$ in \mathcal{P} , if σ occurs exactly once in τ then $[\sigma, \tau]$ is a product of two chains. Moreover, if $\sigma = \tau_{[i,j]}$, then these chains have lengths i and $|\tau| - j + 1$.*

3 Disconnectivity

We say that an interval $[\sigma, \tau]$ is *disconnected* (resp. *shellable*) if $\Delta(\sigma, \tau)$ is disconnected (resp. shellable). Equivalently, $[\sigma, \tau]$ is disconnected if the Hasse diagram of the open interval (σ, τ) is disconnected. For example, in Figure 1.1, the subinterval $[213, 213546]$ is disconnected. In examining the structure of intervals of \mathcal{P} , a natural question is to ask when such intervals $[\sigma, \tau]$ are disconnected, preferably giving the answer in terms of simple conditions on σ and τ . We answer this question in Theorem 3.2 below. Determining when an interval contains a non-trivial disconnected subinterval is a more difficult task, which we explore following Theorem 3.2. Interestingly, as we show in Theorem 3.6, almost all intervals in \mathcal{P} do contain such a disconnected subinterval.

3.1 Characterization of disconnected intervals

It will be helpful to deal with posets of rank 2 separately because we can completely classify them, and because they sometimes require special treatment. In particular, unlike disconnected intervals of rank at least 3, a disconnected interval of rank 2 is shellable; we will study the topic of shellability in the next section. Since every element of \mathcal{P} covers at most 2 elements, a rank-2 interval in \mathcal{P} is either a chain or has two elements of rank 1. Thus the structure of rank-2 intervals is completely determined by Proposition 2.3. With rank-2 disconnected intervals now classified, we will follow [MS15] in saying that a disconnected interval is *non-trivial* if it has rank at least 3.

Definition 3.1 *For $\sigma < \tau$, we say that σ straddles τ if σ is both a prefix and a suffix of τ and has no other occurrences in τ .*

It is easy to check that σ straddles τ if and only if σ is the carrier element of $[\sigma, \tau]$, as defined in the last paragraph of Subsection 2.2. Either way we state it, this is exactly the condition that causes non-trivial disconnectedness of $[\sigma, \tau]$, as we now show.

Theorem 3.2 For $\sigma, \tau \in \mathcal{P}$ with $|\tau| - |\sigma| \geq 3$, we have that $[\sigma, \tau]$ is disconnected if and only if σ straddles τ .

As we will observe in the next section, an interval is non-shellable if it contains a non-trivial disconnected subinterval. Therefore the following direct consequence of Theorem 3.2 is important to the classification of shellable intervals.

Corollary 3.3 An interval $[\sigma, \tau]$ contains a non-trivial disconnected subinterval if and only if there exists $\pi \in [\sigma, \tau]$ such that there are two adjacent occurrences of π in τ that are offset from each other by at least 3 positions. Specifically, if $\tau_i \dots \tau_j$ is the minimal consecutive subsequence of τ containing these two adjacent occurrences, then $[\pi, \tau_{[i,j]}]$ is disconnected.

For example, $[1, 2143576]$ is not itself disconnected, but the occurrences of $\pi = 21$ in 2143576 start at positions 1, 3 and 6, and we can check that $[\pi, \tau_{[3,7]}] = [21, 21354]$ is disconnected.

In the rest of the paper, we will use the term *disconnected subintervals* to mean non-trivial disconnected subintervals.

3.2 Finding disconnected intervals

Finding a disconnected subinterval using Corollary 3.3 may require checking all possible $\pi \in [\sigma, \tau]$. Our next result states that, in some cases, it is sufficient to search in the interval $[x(\tau), \tau]$, which may contain significantly fewer elements than $[\sigma, \tau]$ (as in Example 3.5 below).

Proposition 3.4 Suppose that $|x(\tau)| \neq 2$. If $[\sigma, \tau]$ contains a disconnected subinterval, then so does $[x(\tau), \tau]$.

When $\sigma \leq x(\tau)$, this result nicely complements Theorem 2.1, which says that computing $\mu(\sigma, \tau)$ often boils down to computing $\mu(\sigma, x(\tau))$. So, when determining the Möbius function, the interesting part of the interval is often the part below $x(\tau)$, but when looking for disconnected subintervals, the interesting part is often above $x(\tau)$.

Example 3.5 Consider the interval $[\sigma, \tau] = [1, 68372514]$. We see that $x(\tau) = 2413$, so by Proposition 3.4, we only need to look at $[2413, 68372514]$ when searching for disconnected intervals in $[\sigma, \tau]$. Using Corollary 3.3 and its notation, the only non-trivial candidates for π are 2413 , 35241 and 52413 . In each case, the adjacent occurrences of π in τ are offset by just two positions, so $[\sigma, \tau]$ has no disconnected subintervals. To use Corollary 3.3 without applying Proposition 3.4, we would have had eight non-trivial candidates for π to check.

The condition $|x(\tau)| \neq 2$ is used in the proof of Proposition 3.4 to argue that all adjacent occurrences of $x(\tau)$ are offset by two positions. To see an example of how the statement may fail when $|x(\tau)| = 2$, take $\tau = 1325746$. Its exterior is $x(\tau) = 12$, but the only disconnected subinterval in $[1, \tau]$ is $[21, 21453]$, which is not contained in $[x(\tau), \tau]$.

An immediate consequence of Proposition 3.4 is that $[1, \tau]$ has no (non-trivial) disconnected subintervals if $|\tau| - |x(\tau)| \leq 2$. In fact, we can precisely describe the structure of any interval $[\sigma, \tau]$ where τ has this property, as follows.

When $|\tau| - |x(\tau)| = 1$, Lemma 2.2 implies that τ is monotone, and in this case $[\sigma, \tau]$ is always a chain.

When $|\tau| - |x(\tau)| = 2$, we have that $x(\tau) = \tau_{[1, n-2]} = \tau_{[3, n]}$, where $n = |\tau|$. It follows that

$$\tau_{[i,j]} = \tau_{[i+2, j+2]} \quad \text{for every } 1 \leq i \leq j \leq n - 2. \quad (3.1)$$

Thus, for every $2 \leq k \leq n - 1$, there are exactly two patterns of length k contained in τ , namely $\tau_{[1,k]}$ and $\tau_{[2,k+1]}$. These patterns are different from each other, because otherwise τ would be monotone, implying that $|\tau| - |x(\tau)| = 1$. For any $\sigma \leq \tau$, the interval $[\sigma, \tau]$ has two elements of each length k with $|\sigma| < k < |\tau|$. Additionally, any two elements in $[\sigma, \tau]$ of different lengths are comparable. Note that the two elements of length $n - 2$ are $x(\tau)$ and $i(\tau)$.

We conclude this section with a result that states that, in a certain precise sense, almost all intervals in \mathcal{P} contain disconnected subintervals. The approach below follows that in [MS15], where the analogous result is shown for the classical pattern poset.

Theorem 3.6 *Given a permutation σ , let*

$$\mathbb{P}_n([\sigma, \tau] \text{ contains a disconnected subinterval})$$

denote the probability that $[\sigma, \tau]$ contains a non-trivial disconnected subinterval, where τ is chosen uniformly at random from \mathcal{S}_n . Then

$$\lim_{n \rightarrow \infty} \mathbb{P}_n([\sigma, \tau] \text{ contains a disconnected subinterval}) = 1.$$

The proof of Theorem 3.6 uses the following preliminary result, which is also used in the proof of Corollary 6.9.

Lemma 3.7 *Given a permutation σ , let $\mathbb{P}_n(\sigma \leq \tau)$ denote the probability that $\sigma \leq \tau$, where τ is chosen uniformly at random from \mathcal{S}_n . Then*

$$\lim_{n \rightarrow \infty} \mathbb{P}_n(\sigma \leq \tau) = 1.$$

In contrast with Theorem 3.6, we note that there are infinite classes of intervals that do not contain disconnected subintervals. Given $\sigma \in \mathcal{S}_m$ and $\tau \in \mathcal{S}_n$, we define their *direct sum* $\sigma \oplus \tau$ as the concatenation of σ and the permutation τ^{+m} formed by adding m to every entry of τ . Any time σ satisfies $\sigma_1 > \sigma_{|\sigma|}$, we can take monotone increasing sequences α and β and let $\tau = \alpha \oplus \sigma \oplus \beta$. Then any element of $[\sigma, \tau]$ will occur only once in τ , so there are no disconnected subintervals. For examples of intervals $[\sigma, \tau]$ that are not just a product of two chains but contain no disconnected subintervals, let $\sigma = 21$ and $\tau = \alpha \oplus 21 \oplus 21 \oplus \dots \oplus 21 \oplus \beta$, where the number of copies of 21 is at least 2 and where α and β are monotone increasing as before.

4 Shellability

Our goal for this section is to prove that all intervals in \mathcal{P} are shellable, except for those which are not shellable for a straightforward reason. We begin by giving the necessary background on shellability. We refer the interested reader to [Wac07] for further details and a wealth of other information about poset topology.

When considering a combinatorially defined simplicial complex, such as the order complex $\Delta(\sigma, \tau)$ of an interval $[\sigma, \tau]$ of \mathcal{P} , it is common to ask if the simplicial complex is shellable. Since our intervals $[\sigma, \tau]$ are graded, we can restrict our discussion to order complexes that are pure, i.e., all the facets have the same dimension. A *shelling* of a pure d -dimensional simplicial complex Δ is a linear ordering F_1, F_2, \dots, F_s of its facets such that the intersection

$$\left(\bigcup_{j=1}^{k-1} F_j \right) \cap F_k \tag{4.1}$$

is pure and $(d - 1)$ -dimensional for all $k = 2, 3, \dots, s$. A simplicial complex is *shellable* if it has a shelling. Figure 1.1(b) is an example of a non-shellable complex since in any ordering of the facets, the point 213 will be a $(d - 2)$ -dimensional connected component of an intersection of the type in (4.1), implying that no shelling exists. As is customary, we will say that $[\sigma, \tau]$ is *shellable* if $\Delta(\sigma, \tau)$ is.

In the consecutive pattern poset, it is easy to see that disconnected intervals of rank at least 3 are not shellable. Moreover, a result of Björner [Bjö80] states that if a poset is shellable, then all of its subintervals are shellable. In particular, if an interval $[\sigma, \tau]$ contains a disconnected subinterval of rank at least 3, then it is certainly not shellable. As an example, the disconnected subinterval $[213, 213546]$ in Figure 1.1(a) results in the non-shellability of the corresponding order complex.

It thus follows immediately from Theorem 3.6 that almost all intervals in \mathcal{P} are not shellable, in a particular precise sense.

Corollary 4.1 *Given a permutation σ , let τ be chosen uniformly at random from \mathcal{S}_n . Then the probability that $[\sigma, \tau]$ is shellable tends to 0 as n tends to infinity.*

Our main result of this section is that the converse of Björner’s result is true in \mathcal{P} : if an interval $[\sigma, \tau]$ does not contain a disconnected subinterval, then it is shellable. So, roughly speaking, all intervals in the consecutive pattern poset that have any hope of being shellable are in fact shellable. Recall that we have already classified those intervals that contain disconnected subintervals in Corollary 3.3 in terms of the entries of σ and τ , so the theorem we are about to state classifies those σ and τ that yield a shellable interval $[\sigma, \tau]$.

Theorem 4.2 *The interval $[\sigma, \tau]$ in \mathcal{P} is shellable if and only if it contains no non-trivial disconnected subintervals.*

The proof of the “if” direction is done using CL-shellability. The analogue of Theorem 4.2 in the classical pattern poset is actually false ($[123, 3416725]$ is given in [MS15] as a counterexample).

5 All intervals are rank-unimodal and strongly Sperner

Since intervals in \mathcal{P} are ranked, it is natural to ask about the structure of these intervals with regard to the number of elements at each rank. It is clear from Figure 1.1(a) that the intervals are not rank-symmetric in general. It is also the case that the sequence of rank sizes is not log-concave in general. For example, the interval $[1, 1265473]$ has one element of rank 0 and five elements of rank 2 but only two elements of rank 1. However, there are two interesting properties that intervals of \mathcal{P} have: they are rank-unimodal and strongly Sperner. Our goal from this section is to prove this assertion. We note that neither result is known for the classical pattern poset [MS15]. Even though the definitions of rank unimodal and strongly Sperner are quite different, the two results appear together because they both rely heavily on the same injection between adjacent rank levels given in the following lemma.

Lemma 5.1 *In the interval $[\sigma, \tau]$ in \mathcal{P} , let E_r denote the set of elements of rank r for $0 \leq r \leq N$, where $N = |\tau| - |\sigma|$. For $0 \leq r \leq N - 1$, pick any $m_r \leq \min\{|E_r|, N - r\}$. Then for any selection of m_r elements of E_r , there exists an injection f_r from these elements into E_{r+1} with the property that $\pi < f_r(\pi)$ for each such element π .*

For a finite ranked poset Q , the set of elements of rank i will be called a *rank level* of Q , and the cardinality a_i of this rank level will be referred to as the *size* of rank i . Recall that a finite ranked poset of rank N is called *rank-unimodal* if the sequence a_0, a_1, \dots, a_N of rank sizes is unimodal.

Corollary 5.2 Every interval $[\sigma, \tau]$ in \mathcal{P} is rank-unimodal.

Remark 5.3 In the case of the classical pattern poset, it is an open question whether every interval is rank-unimodal [MS15]. It is known to be true for intervals $[\sigma, \tau]$ with $|\tau| \leq 8$.

In the rest of this section we show that all intervals in \mathcal{P} are strongly Sperner. Let us first give the necessary background. Recall that a ranked poset Q is said to be *Sperner* if the largest rank size equals the size of the largest antichain. Since every rank level of Q consists of an antichain, the Sperner property is equivalent to the condition that some rank level of Q is an antichain of maximum size. For any positive integer k , a *k-family* of Q is a union of k antichains. We are now ready for the definition of one of the well-known generalizations of the Sperner property.

Definition 5.4 Let Q be a poset of rank N . For an integer k with $1 \leq k \leq N + 1$, we say that Q is *k-Sperner* if the sum of the sizes of the k largest ranks equals the size of the largest k -family. In addition, Q is said to be *strongly Sperner* if it is k -Sperner for all k .

Our technique for proving the next theorem makes heavy use of the injections f_r from Lemma 5.1 and of a theorem of Griggs [Gri80] giving an equivalent characterization of strongly Sperner posets in terms of disjoint chains intersecting the largest rank levels.

Theorem 5.5 Every interval $[\sigma, \tau]$ in \mathcal{P} is strongly Sperner.

We conclude this section with an open problem. When σ occurs just once in τ , we know from Proposition 2.4 that $[\sigma, \tau]$ is a lattice.

Problem 5.6 Characterize those intervals $[\sigma, \tau]$ in \mathcal{P} that are lattices. As usual, we would prefer our characterization to be in terms of simple conditions on σ and τ .

For example, for $|\tau| = 4$, $[1, \tau]$ is a lattice if and only if either $\tau_{[1,3]}$ or $\tau_{[2,4]}$ (or both) is monotone, since otherwise both will be upper bounds for 12 and 21.

6 The exterior of a permutation

6.1 The length of the exterior

As we have seen in the previous sections, particularly in Theorem 2.1 and Proposition 3.4, the exterior of a permutation plays an important role in the study of \mathcal{P} . From an enumerative perspective, a natural problem is to study the distribution of the statistic *length of the exterior* on permutations. Table 6.1 shows the number of permutations according to this statistic. The leftmost column counts permutations whose exterior has length 1. Permutations with this property are usually called *non-overlapping permutations* in the literature. It was shown by Bóna [Bón11] that the number of non-overlapping permutations of length n is approximately $0.364n!$, but no exact formula is known. We can show that for every $n \geq 3$, the number of non-overlapping permutations in \mathcal{S}_n is divisible by 4.

The entries in the rightmost nonzero diagonal of Table 6.1 count permutations τ with $|x(\tau)| = |\tau| - 1$. It is clear that the number of such permutations in \mathcal{S}_n is 2 for every $n \geq 2$, since by Lemma 2.2, only the monotone permutations $12 \dots n$ and $n \dots 21$ satisfy this condition. The entries in the diagonal immediately below these 2s count permutations τ with $|x(\tau)| = |\tau| - 2$, which we can enumerate as follows.

$n \setminus k$	1	2	3	4	5	6	7	8	9
2	2								
3	4	2							
4	12	10	2						
5	48	58	12	2					
6	280	306	118	14	2				
7	1864	2186	822	150	16	2			
8	14840	17034	6580	1660	186	18	2		
9	132276	154162	58854	15118	2222	226	20	2	
10	1323504	1532574	588898	150388	30238	2904	270	22	2

Tab. 6.1: The number of permutations $\tau \in \mathcal{S}_n$ with $|x(\tau)| = k$.

Lemma 6.1 For every $n \geq 4$, $|\{\tau \in \mathcal{S}_n : |x(\tau)| = n - 2\}| = 2n + 2$.

Problem 6.2 Find a formula for the other entries of Table 6.1.

6.2 Asymptotic behavior of the length of the exterior

Even without having exact formulas, we can study the behavior of $|x(\tau)|$ as the length of τ goes to infinity. For given n , we will write \mathbb{P}_n and \mathbb{E}_n to denote the probability and the expectation of events involving $\tau \in \mathcal{S}_n$ chosen uniformly at random. In other words, τ is a random variable with uniform distribution over \mathcal{S}_n . We also use the notation $g(n) = o(f(n))$ to mean that $\lim_{n \rightarrow \infty} g(n)/f(n) = 0$.

The main result in this subsection is a constant asymptotic upper bound on the expected length of the exterior of a permutation, given in Theorem 6.4. Its proof relies of the following lemma that bounds the probability that the length of the exterior is large. A similar result appears in [Per13, Lemma 18].

Lemma 6.3 For $1 \leq m \leq n - 1$,

$$\mathbb{P}_n(|x(\tau)| \geq m) \leq \sum_{i=m}^{\lfloor n/2 \rfloor} \frac{1}{i!} + o(n^{-1}),$$

where the summation is defined to be zero when $m > \lfloor n/2 \rfloor$.

Theorem 6.4 The limit of $\mathbb{E}_n(|x(\tau)|)$ as n goes to infinity exists and

$$e - 1 \leq \lim_{n \rightarrow \infty} \mathbb{E}_n(|x(\tau)|) \leq e.$$

Problem 6.5 Find the exact value of $\lim_{n \rightarrow \infty} \mathbb{E}_n(|x(\tau)|)$.

Problem 6.6 Find the limiting distribution of $|x(\tau)|$, that is, find

$$\lim_{n \rightarrow \infty} \mathbb{P}_n(|x(\tau)| = k)$$

for all k .

It was shown by Bóna [Bón11] that when $k = 1$, the sequence $\mathbb{P}_n(|x(\tau)| = 1)$ is decreasing and so its limit exists, and that $0.3640981 \leq \lim_{n \rightarrow \infty} \mathbb{P}_n(|x(\tau)| = 1) \leq 0.3640993$. Even though for fixed k the sequences $\mathbb{P}_n(|x(\tau)| = k)$ are not monotone in general, it is possible to obtain recurrences similar to those in [Bón11] to show that the corresponding limits exist for all k and are positive.

6.3 Permutations with no carrier, and the Möbius function of most intervals.

Recall that τ is said to have a carrier element if $x(\tau) \not\leq i(\tau)$. Carrier elements, defined in [BFS11], play a crucial role in determining the Möbius function of intervals in \mathcal{P} , as we pointed out in Subsection 2.2. A question that arises is how many permutations have a carrier element. In this subsection we study this question, and we use our findings to deduce that most intervals (in a precise sense that we will describe) have a zero Möbius function.

For n from 2 to 10, the number of permutations $\tau \in \mathcal{S}_n$ with no carrier element, that is, satisfying $x(\tau) \leq i(\tau)$, is given by the sequence 0, 4, 12, 84, 548, 4172, 33984, 315800, 3213032,

Problem 6.7 Find a formula for the number of permutations with no carrier element.

Comparing the above sequence with $n!$ suggests that, as n grows, most permutations in \mathcal{S}_n have no carrier element. We can show that this is indeed the case.

Theorem 6.8

$$\lim_{n \rightarrow \infty} \mathbb{P}_n(x(\tau) \not\leq i(\tau)) = 0.$$

A consequence of Theorem 6.8 is that for any fixed σ , most intervals of the form $[\sigma, \tau]$ will have zero Möbius function.

Corollary 6.9 For fixed σ , let $\mathbb{P}_n^\sigma(\mu(\sigma, \tau) = 0)$ denote the probability that $\mu(\sigma, \tau) = 0$ when τ is chosen uniformly at random among permutations in \mathcal{S}_n that contain σ . Then, for every σ ,

$$\lim_{n \rightarrow \infty} \mathbb{P}_n^\sigma(\mu(\sigma, \tau) = 0) = 1.$$

References

- [AEK08] José María Amigó, Sergi Elizalde, and Matthew B. Kennel. Forbidden patterns and shift systems. *J. Combin. Theory Ser. A*, 115(3):485–504, 2008.
- [BFS11] Antonio Bernini, Luca Ferrari, and Einar Steingrímsson. The Möbius function of the consecutive pattern poset. *Electron. J. Combin.*, 18(1):Paper 146, 12 pp., 2011.
- [BJS11] Alexander Burstein, Vít Jelínek, Eva Jelínková, and Einar Steingrímsson. The Möbius function of separable and decomposable permutations. *J. Combin. Theory Ser. A*, 118(8):2346–2364, 2011.
- [Bjö80] Anders Björner. Shellable and Cohen-Macaulay partially ordered sets. *Trans. Amer. Math. Soc.*, 260(1):159–183, 1980.
- [Bón11] Miklós Bóna. Non-overlapping permutation patterns. *Pure Math. Appl. (P.U.M.A.)*, 22(2):99–105, 2011.
- [Bón12] Miklós Bóna. *Combinatorics of permutations*. Discrete Mathematics and its Applications (Boca Raton). CRC Press, Boca Raton, FL, second edition, 2012.
- [Eli15] Sergi Elizalde. A survey of consecutive patterns in permutations. *IMA Volume, Recent Trends in Combinatorics*, 2015. To appear, arXiv:1504.07265.

- [EM15] Sergi Elizalde and Peter McNamara. The structure of the consecutive pattern poset. 2015. preprint, arXiv:1508.05963.
- [EN03] Sergi Elizalde and Marc Noy. Consecutive patterns in permutations. *Adv. in Appl. Math.*, 30(1-2):110–125, 2003. Formal power series and algebraic combinatorics (Scottsdale, AZ, 2001).
- [Gri80] Jerrold R. Griggs. On chains and Sperner k -families in ranked posets. *J. Combin. Theory Ser. A*, 28(2):156–168, 1980.
- [Kit11] Sergey Kitaev. *Patterns in permutations and words*. Monographs in Theoretical Computer Science. An EATCS Series. Springer, Heidelberg, 2011. With a foreword by Jeffrey B. Remmel.
- [MS15] Peter R. W. McNamara and Einar Steingrímsson. On the topology of the permutation pattern poset. *J. Combin. Theory Ser. A*, 134:1–35, 2015.
- [Per13] Guillem Perarnau. A probabilistic approach to consecutive pattern avoiding in permutations. *J. Combin. Theory Ser. A*, 120(5):998–1011, 2013.
- [Smi14] Jason P. Smith. Intervals of permutations with a fixed number of descents are shellable. arXiv:1405.2560 [math.CO], 2014.
- [Smi15] Jason P. Smith. A formula for the Möbius function of the permutation poset based on a topological decomposition. arXiv:1506.04406 [math.CO], 2015.
- [ST10] Einar Steingrímsson and Bridget Eileen Tenner. The Möbius function of the permutation pattern poset. *J. Comb.*, 1(1, [ISSN 1097-959X on cover]):39–52, 2010.
- [Ste13] Einar Steingrímsson. Some open problems on permutation patterns. In *Surveys in combinatorics 2013*, volume 409 of *London Math. Soc. Lecture Note Ser.*, pages 239–263. Cambridge Univ. Press, Cambridge, 2013.
- [SV06] Bruce E. Sagan and Vincent Vatter. The Möbius function of a composition poset. *J. Algebraic Combin.*, 24(2):117–136, 2006.
- [SW12] Bruce E. Sagan and Robert Willenbring. Discrete Morse theory and the consecutive pattern poset. *J. Algebraic Combin.*, 36(4):501–514, 2012.
- [Wac07] Michelle L. Wachs. Poset topology: tools and applications. In *Geometric combinatorics*, volume 13 of *IAS/Park City Math. Ser.*, pages 497–615. Amer. Math. Soc., Providence, RI, 2007.
- [Wil02] Herbert S. Wilf. The patterns of permutations. *Discrete Math.*, 257(2-3):575–583, 2002. Kleitman and combinatorics: a celebration (Cambridge, MA, 1999).