

Decomposing random permutations into order-isomorphic subpermutations

Carla Groenland* Tom Johnston[†] Dániel Korándi^{‡§}
Alexander Roberts[‡] Alex Scott^{‡¶} Jane Tan[‡]

Abstract

Two permutations σ and π are ℓ -similar if they can be decomposed into subpermutations $\sigma^{(1)}, \dots, \sigma^{(\ell)}$ and $\pi^{(1)}, \dots, \pi^{(\ell)}$ such that $\sigma^{(i)}$ is order-isomorphic to $\pi^{(i)}$ for all $i \in [\ell]$. Recently, Dudek, Grytczuk and Ruciński posed the problem of determining the minimum ℓ for which two permutations chosen independently and uniformly at random are ℓ -similar. We show that two such permutations are $O(n^{1/3} \log^{11/6}(n))$ -similar with high probability, which is tight up to a polylogarithmic factor. Our result also generalises to simultaneous decompositions of multiple permutations.

1 Introduction

For two graphs G and H with the same number of vertices and edges, a *U-decomposition* of order ℓ is an edge partition of G and H into edge-disjoint subgraphs G_1, \dots, G_ℓ and H_1, \dots, H_ℓ respectively such that $G_i \cong H_i$ for all $i \in [\ell]$. If such a decomposition exists, we say that G and H are ℓ -similar. Ulam posed the question of determining the minimum $\ell(n)$ such that any two graphs with n vertices and the same number of edges are ℓ -similar (see [16]). This was solved by Chung, Erdős, Graham, Ulam and Yao [8] who showed that the answer is $2n/3 + o(n)$. Chung, Erdős and Graham [6] later also studied the problem of simultaneously decomposing $k \geq 3$ graphs. The answer in this case is $3n/4 + o_k(n)$, and in particular the leading term is independent of k . Surprisingly, it turns out that this bound holds (with $o_k(n)$ replaced by $O(1)$) even if we wish to simultaneously decompose every graph on n vertices with the same number of edges [7]. We refer the reader to [16] for an interesting

*Utrecht University, Utrecht, The Netherlands. Partially supported by CRACKNP with funding from the European Research Council under the EU Horizon 2020 research and innovation programme (grant agreement no. 853234).

[†]School of Mathematics, University of Bristol, Bristol, BS8 1UG, UK and Heilbronn Institute for Mathematical Research, Bristol, UK.

[‡]Mathematical Institute, University of Oxford, Oxford, OX2 6GG, UK.

[§]Supported by SNSF Postdoc.Mobility Fellowship P400P2_186686

[¶]Supported by EPSRC grant EP/V007327/1.

history of this problem, and mention that there are numerous variants which are actively studied (see, for instance, [2, 5, 10]).

The analogous problem for permutations was raised by Dudek, Grytczuk and Ruciński [13]. Given a sequence $\mathbf{a} = a_1, \dots, a_n$ of n distinct real numbers, one can uniquely associate a permutation $\sigma \in S_n$ (which we write as a sequence $\sigma(1), \sigma(2), \dots, \sigma(n)$ with elements in $[n]$) to \mathbf{a} by replacing the i th smallest element of \mathbf{a} with i . We call this σ the *pattern* of \mathbf{a} . Two sequences $\mathbf{a} = a_1, \dots, a_n$ and $\mathbf{b} = b_1, \dots, b_n$ of distinct real numbers are *order-isomorphic* if they have the same pattern, and in this case we write $\mathbf{a} \sim \mathbf{b}$.

Following Dudek, Grytczuk and Ruciński [13], we say that two permutations σ and π are ℓ -similar if they can each be decomposed into ℓ subpermutations $\sigma^{(1)}, \dots, \sigma^{(\ell)}$ and $\pi^{(1)}, \dots, \pi^{(\ell)}$ such that $\sigma^{(i)} \sim \pi^{(i)}$ for each $i \in [\ell]$. We write $U(\sigma, \pi)$ for the smallest ℓ such that σ and π are ℓ -similar. In the worst case, two permutations might only be n -similar. Indeed, the identity permutation $1, \dots, n$ and the reverse permutation $n, \dots, 1$ contain no non-trivial order-isomorphic subpermutations, and so cannot be $(n-1)$ -similar. However, it is far from obvious what happens on average.

Problem 1.1 (Dudek, Grytczuk and Ruciński [13]). *Let σ and π be two permutations of length n chosen independently and uniformly at random. What is the expected value of $U(\sigma, \pi)$?*

A simple $O(\sqrt{n})$ upper bound can be proved by decomposing the permutations into increasing subpermutations as follows. The length of a longest increasing sequence in a uniformly random permutation is asymptotically $2\sqrt{n}$ (see [21]), and this is well-concentrated [3, 14]. In particular, if L is the length of a longest increasing subpermutation in a uniformly random permutation, then

$$\mathbb{P}\left(|L - \mathbb{E}[L]| \geq n^{\frac{1}{3}}\right) \leq \exp(-n^\beta)$$

for some $\beta > 0$. Hence, rather crudely, we have $L \leq 3\sqrt{n}$ with very high probability. By Dilworth's Theorem, this implies that a uniformly random permutation can be decomposed into at most $3\sqrt{n}$ decreasing subpermutations with very high probability. It follows easily that, with very high probability, a pair of uniformly random permutations can be simultaneously decomposed into at most $6\sqrt{n}$ decreasing subsequences of the same length.

Our main result improves the preceding simple bound to the following.

Theorem 1.2. *Let σ and π be two permutations of length n chosen independently and uniformly at random. Then σ and π are $O(n^{1/3} \log^{11/6}(n))$ -similar with probability at least $1 - o(n^{-1})$.*

This immediately provides the following bound on $\mathbb{E}[U(\sigma, \pi)]$.

Corollary 1.3. *Let σ and π be two independent, uniformly random permutations of length n . Then*

$$\mathbb{E}[U(\sigma, \pi)] = O\left(n^{1/3} \log^{11/6}(n)\right).$$

The bound in Theorem 1.2 (and Corollary 1.3) is tight up to the $\log^{11/6}(n)$ factor. Before seeing this, we make a brief detour to survey some closely related work on twinned subpermutations of a single permutation, from which our simple lower bound emerges. Formally, *twins* $T = (T_1, T_2)$ of length ℓ in $\sigma \in S_n$ are a pair of disjoint subsequences $T_1 = \sigma_{i_1}, \dots, \sigma_{i_\ell}$ and $T_2 = \sigma_{j_1}, \dots, \sigma_{j_\ell}$ such that $T_1 \sim T_2$. The problem of finding long twins was introduced by Gawron [15], who showed that the longest twins guaranteed in a permutation of length n is at most $O(n^{2/3})$ and conjectured that this should be tight (up to the constant). As observed in [15], the Erdős-Szekeres theorem provides a lower bound of $\Omega(n^{1/2})$. This lower bound was improved to $\Omega(n^{3/5})$ by Bukh and Rudenko [4].

The proof of Gawron's upper bound comes from applying the first moment method to a uniformly random permutation, and there has since been work on finding a matching lower bound for random permutations. A lower bound of $\Omega(n^{2/3}/\log^{1/3}(n))$ was shown by Dudek, Grytczuk and Ruciński [13], and this was improved to the sharp bound of $\Omega(n^{2/3})$ by Bukh and Rudenko [4]. There has also been further work of Dudek, Grytczuk and Ruciński in the same vein concerning *k-twins* [11], which are a collection of k pairwise disjoint order-isomorphic subpermutations of a single permutation, and other notions of twins with additional restrictions imposed or a weaker similarity condition (see [9, 12]).

Returning to Problem 1.1, if one chooses a permutation σ of length $2n$ uniformly at random, then the patterns of the first and the second half of σ are independent, uniformly random permutations of length n . Order-isomorphic subpermutations of length ℓ between the first and second of half of σ give rise to twins of length ℓ in σ . Hence, from Gawron's result that the longest twin in a permutation of length $2n$ is of length $O(n^{2/3})$ with high probability [15], it follows that the longest order-isomorphic subpermutations of two uniformly random permutations of length n are also of length $O(n^{2/3})$ with high probability. This means that $\Omega(n^{1/3})$ pieces are needed with high probability.

As with graphs, the definition of similarity generalises naturally to multiple permutations. We say that a collection of k permutations $\sigma_1, \dots, \sigma_k$ are ℓ -similar if each permutation σ_j can be split into ℓ subpermutations $\sigma_j^{(1)}, \dots, \sigma_j^{(\ell)}$ such that $\sigma_j^{(i)} \sim \sigma_{j'}^{(i)}$ for all choices of $i \in [\ell]$, and $j, j' \in [k]$. Mirroring the progression of work on U-decompositions, this leads to the problem of determining, for each k , the minimum $\ell(n)$ such that a set of k permutations of length n chosen independently and uniformly at random are ℓ -similar. There is again a simple upper bound of $O(\sqrt{n})$ which comes from decomposing the permutations into increasing subpermutations. Using the result of Dudek, Grytczuk and Ruciński [11] which states that the longest k -twins in a random permutation of length n have length at most $O(n^{k/(2k-1)})$ with high probability, we can deduce a lower bound of $\Omega(n^{(k-1)/(2k-1)})$ in the same manner as the $k = 2$ case.

Our proof of Theorem 1.2 extends to give a bound in the setting of multiple permutations, and it is again tight up to a polylogarithmic factor.

Theorem 1.4. *For any fixed integer $k \geq 2$, let $\sigma_1, \dots, \sigma_k$ be permutations of length n chosen independently and uniformly at random. Then $\sigma_1, \dots, \sigma_k$ are*

$O\left(n^{\frac{k-1}{2k-1}} \log^{\frac{3(k-1)}{2} + \frac{1}{2k-1}}(n)\right)$ -similar with probability at least $1 - o(n^{-1})$.

2 Proof of the main theorems

We begin with an outline of the proof of Theorem 1.2; the minor modifications required to prove Theorem 1.4 are given at the end of this section.

In the following, the idea of constructing an auxiliary bipartite multigraph has previously been used by Dudek, Grytczuk and Ruciński [13] and later by Bukh and Rudenko [4] to study twins in a single random permutation. The use of a Poisson process to sample random permutations also appears in [4].

It is convenient to work with random permutations generated from n uniformly distributed points in the unit square $[0, 1]^2$. With the n points ordered by their x -coordinate, the permutation can be read off as the relative ordering of the y -coordinates. This representation was used by Bukh and Rudenko [4]. Given sets of n “red” points and n “blue” points which define a pair of random permutations, we choose a matching between the red points and the blue points and only consider decompositions into subpermutations where the matched points are corresponding entries in the decomposition.

A useful strategy to simplify viewing patterns of permutations generated in this way is to discretise the unit square by splitting it up into an $M \times M$ grid of smaller squares, and consider the squares in which the points fall. Then, given a subpermutation in which there is at most one point in each row and column, the pattern of the subpermutation is determined by the squares in which each point lies and we do not need to know their positions within the squares. If we wish to use this subpermutation in our decomposition, we need the pattern of this subpermutation to be the same for the red points as for the matched blue points. We only consider the simple case where all of the matched points are offset by the same amount, and so we label each point by the relative position (measured in grid squares) of its matched point. We will handle each label using separate subpermutations, making it necessary to avoid having too many different possible labels. We will achieve this by ensuring matched points are close together. The following result of Leighton and Shor allows us to control the maximum distance between matched points.

Theorem 2.1 (Leighton, Shor [18]). *Suppose A and B are two independent sets of n points uniformly distributed in $[0, 1]^2$. Then there is a function $\alpha = \alpha(n) = \Omega(\sqrt{\log n})$ and an absolute constant C such that, with probability at least $1 - n^{-\alpha}$, there is a perfect matching between A and B for which the maximum distance between two matched points is at most $C \log^{3/4}(n)/\sqrt{n}$.*

To complete the proof, we further partition the points corresponding to a specific label in such a way that each part contains at most one point from each row and column. This can be reduced to the problem of edge-colouring an auxiliary bipartite multigraph, and it suffices to bound the number of points in a given region. If we were to naively generate n points uniformly at random, then the numbers of points in different regions are not independent, even if the

regions are disjoint. Thus, we generate points using a Poisson process with mean $2n$ and randomly delete excess points to obtain n uniformly distributed points. The number of our n points in a given region is bounded by the number of points of the Poisson process in that region, and this is independent for disjoint regions.

We will make use of the following well-known bound on the tail probabilities of the Poisson distribution.

Lemma 2.2 (Theorem 5.4 in [19]). *Let $X \sim \text{Poisson}(\lambda)$ and $x > \lambda$. Then*

$$\mathbb{P}(X \geq x) \leq \frac{(e\lambda)^x e^{-\lambda}}{x^x}.$$

We are now ready to give the formal argument.

Proof of Theorem 1.2. Given n points $\{(a_i, b_i) : i \in [n]\}$ in the unit square $[0, 1]^2$ such that $a_1 < a_2 < \dots < a_n$, we can define a permutation σ by setting $\sigma(i) = j$, where b_i is the j th smallest element of the set $\{b_1, \dots, b_n\}$. If the n points are chosen uniformly at random in the unit square, this process generates a uniformly random permutation.

Let R be a 2D Poisson process with rate $2n$ and let N_R denote the number of points of R in $[0, 1]^2$. Then $N_R \sim \text{Poisson}(2n)$ and

$$\mathbb{P}(N_R \leq n) \leq \exp\left(-\frac{n^2}{2n}\right).$$

Conditional on the number of points in $[0, 1]^2$ being $m \geq n$, the m points are distributed uniformly. After removing $m - n$ points uniformly at random, we have n uniformly distributed points, which we call the *red points*, and we can use this to get a uniformly random permutation $\sigma \in S_n$. Similarly, we can use an independent Poisson process B to define a set of n uniformly distributed *blue points* and a corresponding permutation $\pi \in S_n$.

We now match up the n red points with the n blue points using Theorem 2.1.

Let $M = \lceil n^{2/3} \log^{-1/3}(n) \rceil$ and split the square $[0, 1]^2$ into an $M \times M$ grid with evenly spaced rows and columns. Let $S_{i,j}$ denote the square in the i th row and j th column. That is,

$$S_{i,j} = \left[\frac{j-1}{M}, \frac{j}{M} \right) \times \left[\frac{i-1}{M}, \frac{i}{M} \right).$$

Suppose that we have a set of red points such that no two red points are in the same row or column. Then the pattern of the corresponding subpermutation is entirely determined by the squares, and we do not need to know the particular points. If the squares for the matched blue points follow the same pattern, then the red points and blue points correspond to order-isomorphic subpermutations of σ and π .

To find such sets of points, we will look for matchings in which all edges have the same label in an auxiliary edge-labelled bipartite multigraph G constructed




Figure 1: The square $[0, 1]^2$ split into $M \times M$ boxes. The red and blue points have been sampled from a Poisson point process with mean 50, and 25 of each have been selected uniformly at random to form the red and blue permutations respectively. The selected points have been matched using a minimax matching.

as follows. Let G have vertex set $\{i_1, \dots, i_M\} \cup \{j_1, \dots, j_M\}$. Given a red point r in row r_{row} and column r_{col} matched with a blue point b in row b_{row} and column b_{col} , add an edge from row $i_{r_{\text{row}}}$ to column $j_{r_{\text{col}}}$ in the graph and label it $(b_{\text{row}} - r_{\text{row}}, b_{\text{col}} - r_{\text{col}})$, i.e. label it with the relative position of b from r . Note that a matching in G represents points in different rows and columns, so a matching in which all edges have the same label, as explained above, corresponds to order-isomorphic subpermutations of σ and π . It will therefore suffice to decompose the edges of G into such matchings. Equivalently, we wish to separately properly edge-colour the edges with each label. Note that as each matched edge has length at most $C \log^{3/4}(n)/\sqrt{n}$, there are at most

$$\left(2M \frac{C \log^{3/4}(n)}{\sqrt{n}} + 2\right)^2 = O\left(\frac{M^2 \log^{3/2}(n)}{n}\right) \quad (2.1)$$

possible labels. We will handle each label separately.

It is well-known that a bipartite multigraph can be edge-coloured using Δ colours where Δ is the maximum degree of the graph. In fact, this can be done efficiently (see e.g. [20]). We bound the degree in two steps: first we bound the number of distinct edges labelled (c_1, c_2) that a vertex sees, then we bound the maximum number of repetitions of an edge.

There is an edge from row i to column j with label (c_1, c_2) if and only if there is a red point in the square $S_{i,j}$ matched to a blue point in the square $S_{i+c_1, j+c_2}$. In particular, this means that there must be a point from the Poisson process R in $S_{i,j}$ and from the process B in the square $S_{i+c_1, j+c_2}$. This happens with




Figure 2: A sample minimax matching along with the bipartite graph for the label $(0,0)$. The 5 edges can be coloured using 2 colours, and so the subpermutations of length 5 can be decomposed into 2 parts.

probability

$$\left(1 - \exp\left(-\frac{2n}{M^2}\right)\right)^2 \leq \frac{4n^2}{M^4}.$$

As the number of points of a Poisson process in a particular square is independent of the numbers in the other squares, the number of distinct edges labelled (c_1, c_2) that are incident with a given vertex in G , is dominated by a binomial random variable with M trials and success probability $4n^2/M^4$. Let $X \sim \text{Bin}(M, 4n^2/M^4)$. Then

$$\mathbb{E}[X] = 4n^2/M^3 \leq \frac{4n^2}{\left(n^{2/3} \log^{-1/3}(n)\right)^3} = 4 \log n.$$

We also have

$$4n^2/M^3 \geq \frac{4n^2}{\left(n^{2/3} \log^{-1/3}(n) + 1\right)^3} \geq 3 \log n$$

for $n \geq 64$. Hence, using a Chernoff bound (see e.g. [19]),

$$\begin{aligned} \mathbb{P}(X \geq 12 \log n) &\leq \mathbb{P}(X \geq 12n^2/M^3) \\ &\leq \exp(-4n^2/M^3) \\ &\leq \exp(-3 \log n) \\ &= n^{-3}. \end{aligned}$$

There are $2M$ vertices in G and there are $O(M^2 \log^{6/4}(n)/n)$ labels (see (2.1)), so the probability that any vertex is incident with at least $12 \log n$ distinct

edges of the same label is at most

$$n^{-3} \cdot 2M \cdot O\left(\frac{M^2 \log^{3/2}(n)}{n}\right) = O\left(\frac{\log^{1/2}(n)}{n^2}\right) = o(n^{-1}).$$

We now look to bound the number of times a particular edge appears in the graph. For there to be 8 edges from row i to column j , the square $S_{i,j}$ must contain at least 8 points of the Poisson process R . The number of such points in any given square is a Poisson random variable with mean

$$2n/M^2 \leq \frac{2 \log^{2/3}(n)}{n^{1/3}}.$$

Hence, using the bound from Lemma 2.2, the probability that there are at least 8 points in $S_{i,j}$ is at most

$$\frac{(2e)^8 \log^{16/3}(n)}{8^8 n^{8/3}} \exp\left(-\frac{2 \log^{2/3}(n)}{n^{1/3}}\right) \leq \frac{\log^{16/3}(n)}{n^{8/3}}.$$

Taking the union bound over all M^2 squares, the probability that any square contains 8 points is at most $O(\log^{14/3}(n)/n^{4/3}) = o(n^{-1})$.

Consider the subgraph of G formed by the edges with label (c_1, c_2) . With high probability, every edge is repeated at most 7 times and every vertex sees at most $12 \log n$ distinct edges. Hence, the maximum degree is bounded by $84 \log n$, and so we can cover the edges with $84 \log n$ matchings. We can do this separately for every distinct label and cover all edges in at most

$$O\left(\frac{M^2 \log^{6/4}(n)}{n}\right) \cdot 84 \log n = O\left(n^{1/3} \log^{11/6}(n)\right)$$

matchings, each consisting of edges with the same label as required. \square

We remark that this proof gives a simple, efficient algorithm to decompose two random permutations σ and π into twins. More explicitly, given a permutation σ of length n , generate n red points in $[0, 1]^2$ by first sampling n values x_1, \dots, x_n uniformly in $[0, 1]$ to be the x coordinates and then sampling n values y_1, \dots, y_n uniformly in $[0, 1]$ to be the y coordinates. By relabelling the values as necessary, we can assume that $x_1 < \dots < x_n$ and $y_1 < \dots < y_n$. Let the n red points in $[0, 1]^2$ be $(x_i, y_{\sigma(i)})$, and similarly choose n blue points in $[0, 1]^2$ for π . We then find a minimax matching between the red and blue points. This can easily be done in $O(n^{5/2} \log n)$ time by using a binary search over the maximum allowed edge weight (from the $\binom{n}{2}$ possibilities) and checking if there is a perfect matching using the allowed edges in $O(n^{5/2})$ time [17]. Crudely, there are at most n possible labels, and for each label we need to edge-colour a bipartite multigraph with at most n edges. Each colouring can be done in $O(n \log n)$ time [1, 20]. This process therefore decomposes the permutations into order-isomorphic subpermutations in $O(n^{5/2} \log n)$ time.

The decomposition produced by the above process may not be optimal. However, the proof of Theorem 1.2 shows that the algorithm uses on average $O(n^{1/3} \log^{11/6}(n))$ subpermutations, which is not far from the simple lower bound of $\Omega(n^{1/3})$.

The preceding argument can be generalised with very few modifications to prove Theorem 1.4. By starting with k independent Poisson processes with rate $2n$, we can generate k collections of n uniformly distributed points in $[0, 1]^2$ giving rise to independent uniformly random permutations $\sigma_1, \dots, \sigma_k \in S_n$. Let the n points determining σ_i be P_i . The points P_1 will play the role of the red points in the proof of Theorem 1.2, while the other sets of points mirror the blue points. Accordingly, we match up the points P_1 and P_i for each $i = 2, \dots, k$ using minimax matchings. By Theorem 2.1, with probability $1 - o(n^{-1})$, no edge in any of the matchings is longer than $C \log^{3/4}(n)/\sqrt{n}$.

Take $M = \left\lceil n^{\frac{1}{2} + \frac{1}{2(2k-1)}} \log^{-\frac{1}{2k-1}}(n) \right\rceil$ and split the square $[0, 1]^2$ into an $M \times M$ grid with evenly spaced rows and columns. We construct a similar bipartite multigraph as in the proof of Theorem 1.2. Namely, suppose we have a ‘red’ point p_1 matched to points p_2, \dots, p_k , where $p_x \in P_x$ is in row r_x and column c_x for each $x \in [k]$. We then add an edge from row vertex i_{r_1} to column vertex j_{c_1} in the graph and label it by

$$(r_2 - r_1, c_2 - c_1, r_3 - r_1, c_3 - c_1, \dots, r_k - r_1, c_k - c_1).$$

With probability $1 - o(n^{-1})$, there are at most

$$O\left(\frac{M^{2(k-1)} \log^{3(k-1)/2}(n)}{n^{k-1}}\right)$$

distinct labels. To bound the degree of the multigraph, it suffices to show that the probability of a vertex being incident to $3 \cdot 2^k \log n$ distinct edges with the same label is $o(n^{-1})$, and the probability that an edge has multiplicity at least $4k$ is also $o(n^{-1})$. This can be done as in the proof of Theorem 1.2. Hence, with probability $1 - o(n^{-1})$, the permutations are ℓ -similar for

$$\begin{aligned} \ell &= O\left(\frac{M^{2(k-1)} \log^{3(k-1)/2}(n)}{n^{k-1}}\right) \cdot 3 \cdot 2^k (4k - 1) \log n \\ &= O\left(n^{\frac{k-1}{2k-1}} \log^{\frac{3(k-1)}{2} + \frac{1}{(2k-1)}}(n)\right). \end{aligned}$$

3 Open problems

Theorem 1.2 is tight up to the polylogarithmic factor and it would be interesting to know whether some polylogarithmic factor is necessary, or if two uniformly random permutations are $O(n^{1/3})$ -similar. Even if a polylogarithmic factor is necessary, it is likely that the $\log^{11/6}(n)$ we have shown can be improved. For example, we have used a minimax matching which guarantees that the number

of labels is small and assumed that every label is equally bad, but it is likely that some of the labels have fewer edges and the edges can be covered using fewer matchings. Indeed, if we replace the minimax matching with a matching that minimises the transportation distance, the average distance between a red point and the matched blue point is $\Theta(\log^{1/2}(n)/\sqrt{n})$, saving a factor of $\log^{1/4}(n)$. This suggests that most edges should be spread over fewer labels and that the polylogarithmic factor may be improved upon.

Problem 3.1. *Are two permutations of length n chosen independently and uniformly at random $O(n^{1/3})$ -similar with high probability?*

The other main problem pertaining to (general) twins in permutations is to determine the maximum size of a twin that one is guaranteed to find in a permutation of length n . The best lower bound, due to Bukh and Rudenko [4], is currently $\Omega(n^{3/5})$, while the best upper bound is $O(n^{2/3})$ obtained by applying a first moment calculation to a random permutation [15]. It would be interesting to narrow this gap. Resolving whether the upper bound is tight would be of particular interest, as this would determine whether there are permutations for which the maximum twin length is smaller than that of the random permutation.

Problem 3.2. *Does every permutation of length n contain a twin of length $\Omega(n^{2/3})$?*

References

- [1] N. Alon. A simple algorithm for edge-coloring bipartite multigraphs. *Information Processing Letters*, **85**(6):301–302, 2003.
- [2] M. Axenovich. Repetitions in graphs and sequences. In A. Beveridge, J. R. Griggs, L. Hogben, G. Musiker and P. Tetali, editors, *Recent Trends in Combinatorics*. Springer, 2016.
- [3] B. Bollobás and G. Brightwell. The height of a random partial order: concentration of measure. *The Annals of Applied Probability*, **2**(4):1009–1018, 1992.
- [4] B. Bukh and O. Rudenko. Order-isomorphic twins in permutations. *SIAM Journal on Discrete Mathematics*, **34**(3):1620–1622, 2020.
- [5] B. Bukh and L. Zhou. Twins in words and long common subsequences in permutations. *Israel Journal of Mathematics*, **213**(1):183–209, 2016.
- [6] F. Chung, P. Erdős and R. L. Graham. Minimal decompositions of graphs into mutually isomorphic subgraphs. *Combinatorica*, **1**(1):13–24, 1981.
- [7] F. Chung, P. Erdős and R. L. Graham. Minimal decomposition of all graphs with equinumerous vertices and edges into mutually isomorphic subgraphs. In *Finite and Infinite Sets*, volume 37 of *Colloquia Mathematica Societatis János Bolyai*, pages 171–179. North-Holland, 1984.

- [8] F. Chung, P. Erdős, R. L. Graham, S. Ulam and F. Yao. Minimal decompositions of two graphs into pairwise isomorphic subgraphs. In *Proceedings of the Tenth Southeastern Conference on Combinatorics, Graph Theory and Computing*, pages 3–18, 1979.
- [9] A. Dudek, J. Grytczuk and A. Ruciński. On weak twins and up-and-down sub-permutations. *preprint*, 2020. arXiv:2012.11451.
- [10] A. Dudek, J. Grytczuk and A. Ruciński. Long twins in random words. *preprint*, 2021. arXiv:2112.14197.
- [11] A. Dudek, J. Grytczuk and A. Ruciński. Multiple twins in permutations. *preprint*, 2021. arXiv:2107.06974.
- [12] A. Dudek, J. Grytczuk and A. Ruciński. Tight multiple twins in permutations. *preprint*, 2021. arXiv:2105.02591.
- [13] A. Dudek, J. Grytczuk and A. Ruciński. Variations on twins in permutations. *Electronic Journal of Combinatorics*, **28**(3):P3.19, 2021.
- [14] A. Frieze. On the length of the longest monotone subsequence in a random permutation. *The Annals of Applied Probability*, **1**(2):301–305, 1991.
- [15] M. Gawron. Izomorficzne podstruktury w słowach i permutacjach. Master’s thesis, Uniwersytet Jagielloński, 2014.
- [16] R. Graham. Reflections on a theme of Ulam. In R. Gera, T. W. Haynes and S. T. Hedetniemi, editors, *Graph Theory: Favorite Conjectures and Open Problems - 2*, pages 55–61. Springer, 2018.
- [17] J. E. Hopcroft and R. M. Karp. An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs. *SIAM Journal on Computing*, **2**(4):225–231, 1973.
- [18] T. Leighton and P. Shor. Tight bounds for minimax grid matching with applications to the average case analysis of algorithms. *Combinatorica*, **9**(2):161–187, 1989.
- [19] M. Mitzenmacher and E. Upfal. *Probability and Computing: Randomization and Probabilistic Techniques in Algorithms and Data Analysis*. Cambridge University Press, 2017.
- [20] A. Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*, volume 24 of *Algorithms and Combinatorics*. Springer, 2003.
- [21] J. M. Steele. Variations on the monotone subsequence theme of Erdős and Szekeres. In D. Aldous, P. Diaconis, J. Spencer and J. M. Steele, editors, *Discrete Probability and Algorithms*, pages 111–131. Springer, 1995.