hexagon:2611.01117v1

COUNTING CONTINGENCY TABLES, KOSTKA NUMBERS, AND LITTLEWOOD–RICHARDSON COEFFICIENTS WITH UNARY INPUTS IS #P-COMPLETE

GRETA PANOVA

Abstract. We prove that counting contingency tables with prescribed unary margins is #P-complete under polynomial-time Turing reductions for both cases of tables with nonnegative integer entries and entries in $\{0,1\}$. This proves a Conjecture of Pak and Panova. The methods involve modular arithmetic by repeating margins p-times and relating to the SETPARTITION problem. The congruences are proven using symmetric function techniques. We also realize each nonnegative table count as a single rectangular Kostka number, and then as a Littlewood–Richardson coefficient with a rectangular factor. Both coefficient problems are therefore #P-complete for partitions given in unary making Kostka and Littlewood-Richardson coefficients strongly #P-complete. For permutations in one-line notation, computing a Schubert structure constant is consequently #P-hard even when all three permutations are Grassmannian with a common descent.

1. Introduction

Let $\mathrm{T}(a,b)$ be the number of contingency tables with row sums $a=(a_1,\dots,a_n)$ and column sums $b=(b_1,\dots,b_m)$. Dyer, Kannan, and Mount proved that computing $\mathrm{T}(a,b)$ is #P-complete when the margins are input in binary (so input size is $O(\sum_i \log_2(a_i)+\sum_j \log_2(b_j))$), even when $n=2$ [DKM97, Thm. 1]. The distinction between binary and unary input already appears in this two-row case. If $a=(r,N-r)$, the first row determines the table, whence

$$\mathrm{T}(a,b)=[z^r]\prod_{j=1}^{m}(1+z+\cdots+z^{b_j}).$$

Truncating each multiplication at degree r gives an exact algorithm polynomial in m and N. This is polynomial time for unary margins.

Miller and Harrison [MH13] give polynomial-time exact algorithms for both nonnegative and binary tables when the largest column sum is bounded by a constant. They also raise the hardness question for $\{0,1\}$-tables. Pak and Panova conjectured that when the inputs are encoded in unary the computational problem is still #P-complete for nonnegative tables, Kostka numbers, and Littlewood–Richardson coefficients [Pak24, §13.4, items (1)–(3)], see also [Pan25]. Pak and Robichaux again describe the table problem as conjectural in 2025 [PR25b, p. 4]. Narayanan had proved hardness of both tableau coefficients with binary input [Nar06]. Recent manuscripts of OpenAI give an exact uniform sampler for nonnegative tables [Ope26a] and a randomized approximation scheme for counting tables with cell bounds [Ope26b].

Write $\mathrm{B}(a,b)$ for the number of contingency tables with entries in $\{0,1\}$, which are sometimes confusingly called *binary contingency tables*. The input in both cases consists solely of the two lists of margins.

Theorem 1.1. *Computing* $\mathrm{T}(a,b)$ *and* $\mathrm{B}(a,b)$ *is* #P*-complete under deterministic polynomial-time Turing reductions when the margins* $a=(a_1,\dots,a_n)$ *and* $b=(b_1,\dots,b_n)$ *are given in unary.*

A binary table is the incidence matrix between the two labeled vertex classes of a simple bipartite graph. Thus Theorem 1.1 also proves #P-completeness for counting labeled realizations of a prescribed bipartite degree sequence.

The hardness comes from the 0–1 permanent. We encode a perfect matching by assigning indivisible items to labeled bins. Repeating each margin a prime number p of times lets us recover the assignment count modulo p from one exact table count. A row of sum d has generating function h_d for nonnegative tables and e_d for binary tables. A single specialization argument applies to both. The Chinese remainder theorem then recovers the integer assignment count.

Here in Sections 5 and 6 we show that computing Kostka numbers $K_{\lambda,\mu}$ and Littlewood-Richardson coefficients $c^{\lambda}_{\mu\nu}$ is #P-complete when the inputs are in unary, namely λ, μ, ν are given as Young diagrams and the input size is $O(|\lambda|)$. This also shows, see 7, that computing Schubert structure constants C^w_{uv} is #P-hard where the input is given by the permutations w, u, v, a question in [Pak24]. This follows since $C^w_{uv} = c^{\lambda(w)}_{\lambda(u)\lambda(v)}$ when u, v, w are Grassmannian with unique descent at k, and $\lambda(w) := (w_k - k, w_{k-1} - k + 1, \ldots, w_1 - 1)$.

2. Tables and ordered set partitions

Write $[k] = \{1, \ldots, k\}$. For $a \in \mathbb{Z}^n_{\geq 0}$ and $b \in \mathbb{Z}^m_{\geq 0}$ of common total N, set

$$\mathcal{T}(a,b) = \left\{ X \in \mathbb{Z}^{n\times m}_{\geq 0} : \sum_j X_{ij} = a_i,\ \sum_i X_{ij} = b_j \right\}, \qquad \mathrm{T}(a,b) = |\mathcal{T}(a,b)|.$$

$\mathrm{B}(a,b)$ is defined analogously with the restriction that $X \in \{0,1\}^{n\times m}$. Rows and columns are labeled, including those with equal margins. Unequal totals give zero. Zero margins can be deleted from a table instance, and total zero gives a single table. We therefore use positive margins below. A unary list occupies $O(N + n + m)$ symbols, including separators.

For the same positive vectors define the *ordered set partition count*

$$\mathrm{P}(a,b) = \#\left\{ f : [n] \longrightarrow [m] : \sum_{i:\, f(i)=j} a_i = b_j \text{ for all } j \right\}. \tag{1}$$

This is the ordered set partition problem of Ikenmeyer, Pak, and Panova [IPP24, §3.3]. The indices of the items distinguish them even if their weights agree. For example, $\mathrm{T}((2,2),(1,3)) = 2$, whereas $\mathrm{P}((2,2),(1,3)) = 0$ since neither bin can receive an odd total.

For an integer $p \geq 1$, write $a^{[p]}$ for the vector of length pn containing p successive copies of each a_i. Thus $(2,3)^{[2]} = (2,2,3,3)$. We use $F \leq^{\mathrm{p}}_{\mathrm{T}} G$ for a deterministic polynomial-time oracle reduction of a function F to a function G. Such a reduction is *nonadaptive* if the query instances can all be constructed before receiving any answers.

3. A congruence for repeated margins

Theorem 3.1. *Let a, b be positive margin vectors of common total N, with n row margins. Let p be a prime larger than every a_i and b_j. Put*

$$T_p = \mathrm{T}(a^{[p]}, b^{[p]}), \qquad B_p = \mathrm{B}(a^{[p]}, b^{[p]}), \qquad A = \prod_{i=1}^{n} a_i.$$

Both counts are divisible by p^n, and

$$(-1)^n A \frac{T_p}{p^n} \ \equiv\ \mathrm{P}(a,b) \ \equiv\ (-1)^N A \frac{B_p}{p^n} \pmod{p}. \tag{2}$$

Let h_d and e_d be the complete homogeneous and elementary symmetric functions. In a finite alphabet, h_d is the sum of all monomials of degree d, and e_d is the sum of the squarefree ones. For $y = (y_1, \ldots, y_m)$, write $y^{[p]}$ for the alphabet with p copies of each y_j. The following specialization accounts for the two signs in (2).

Lemma 3.2. *For $1 \le d < p$, both $h_d(y^{[p]})$ and $e_d(y^{[p]})$ are divisible by p coefficientwise, and*

$$\frac{h_d(y^{[p]})}{p} \equiv \frac{1}{d}\sum_{j=1}^{m} y_j^d, \qquad \frac{e_d(y^{[p]})}{p} \equiv \frac{(-1)^{d-1}}{d}\sum_{j=1}^{m} y_j^d \pmod p.$$

Proof. For $|\alpha| = d$, the coefficient of y^α in either specialization is $\prod_j c(\alpha_j)$, with $c(r) = \binom{p+r-1}{r}$ for h_d ($y_i^{\alpha_i}$ appears from a weak composition $(\alpha_{i1}, \alpha_{i2}, \dots, \alpha_{ip})$ of α_i from the monomial $y_i^{\alpha_{i1}} y_i^{\alpha_{i2}} \cdots y_i^{\alpha_{ip}}$) and simiarly $c(r) = \binom{p}{r}$ for e_d. When $1 \le r < p$,

$$\frac{r}{p}\binom{p+r-1}{r} = \binom{p+r-1}{r-1} \equiv 1, \qquad \frac{r}{p}\binom{p}{r} = \binom{p-1}{r-1} \equiv (-1)^{r-1} \pmod p.$$

Every positive α_j therefore contributes a factor of p. After division by p, only the pure powers y_j^d survive modulo p, with the displayed coefficients. □

Proof of Theorem 3.1. Write $g = h$ or $g = e$, and set $C_p = T_p$ or $C_p = B_p$, respectively. Introduce a variable $x_{j,t}$ for each replicated column, $1 \le j \le m$ and $1 \le t \le p$, and put $\beta_{j,t} = b_j$. The row generating function gives

$$C_p = [x^\beta] \prod_{i=1}^{n} g_{a_i}(x)^p.$$

Set $x^p = (x_{j,t}^p)_{j,t}$. By the Frobenius congruence,

$$W_d(x) = \frac{g_d(x)^p - g_d(x^p)}{p}$$

is an integer polynomial. Since $a_i > 0$, every monomial of $g_{a_i}(x^p)$ has an exponent at least p. All entries of β are smaller than p, so in extracting $[x^\beta]$ the terms from $g_{a_i}(x^p)$ do not contribute and we may replace each $g_{a_i}(x)^p$ by $pW_{a_i}(x)$. Thus p^n divides C_p and

$$C_p/p^n = [x^\beta] \prod_i W_{a_i}(x). \tag{3}$$

Now let $F(x)$ be any integer polynomial symmetric in each block of p variables. Identifying the variables of block j with y_j gives

$$[x^\beta]F(x) \equiv [y^{pb}]F(y^{[p]}) \pmod p. \tag{4}$$

Indeed, group monomials with the same multiset of exponents in each block. If exponent r occurs k_r times in a block, the number of arrangements in that block is $p!/\prod_r k_r!$. This is divisible by p unless all p exponents agree. Symmetry gives equal coefficients for all arrangements, so a contribution survives modulo p only when the exponents in each block are constant. For block sums pb_j, the unique such monomial is x^β.

Let $\varepsilon_h(d) = 1$ and $\varepsilon_e(d) = (-1)^{d-1}$. By Lemma 3.2, $g_d(y^{[p]})^p/p$ vanishes modulo p. Using the same lemma with y_j^p in place of y_j in the second term of W_d gives

$$W_d(y^{[p]}) \equiv -\frac{\varepsilon_g(d)}{d}\sum_j y_j^{pd} \pmod p.$$

Equations (3) and (4) now give

$$\frac{C_p}{p^n} \equiv \frac{(-1)^n \prod_i \varepsilon_g(a_i)}{A}[y^{pb}] \prod_i \left(\sum_j y_j^{pa_i}\right) = \frac{(-1)^n \prod_i \varepsilon_g(a_i)}{A}\ \mathrm{P}(a,b) \pmod p.$$

Choosing a term for item i assigns its full weight to bin j. The sign is $(-1)^n$ for h and $(-1)^{n+\sum_i (a_i-1)} = (-1)^N$ for e. □

Example 3.3. For $a = (1, 1, 2)$ and $b = (2, 2)$ there are two assignments, according to which bin receives the item of weight 2. For $p = 3$ the replicated table has nine rows and six columns. In six variables, extracting $[x_1^2 \cdots x_6^2]$ from $h_1^6 h_2^3$ and $e_1^6 e_2^3$ gives $T_3 = 1{,}220{,}400$ and $B_3 = 707{,}400$. The signs in (2) differ. Both normalized counts recover the two assignments, since

$$-2 \cdot 45{,}200 \equiv 2 \cdot 26{,}200 \equiv 2 \pmod 3.$$

4. The complexity reduction

For either choice of table oracle the reduction is

$$\mathrm{per}_{0,1} \ \leq_{\mathrm{T}}^{\mathrm{p}} \ \mathsf{SetPartition}_{\mathrm{unary}} \ \leq_{\mathrm{T}}^{\mathrm{p}} \ \mathcal{C}, \qquad \mathcal{C} \in \{\mathsf{ContingencyTable}_{\mathrm{unary}}, \mathsf{BinaryTable}_{\mathrm{unary}}\}. \quad (5)$$

The first step is included to make the source of hardness explicit. A related construction for higher-dimensional matchings appears in [IPP24, §4.3].

Proposition 4.1. *The* 0–1 *permanent reduces to counting ordered set partitions of unary weights. A single assignment count suffices.*

Proof. View the matrix as the adjacency matrix of a bipartite graph $G = (L, R; E)$ with q vertices in each class, indexed by $[q]$. Its permanent counts perfect matchings. The cases $q = 0$ or of an isolated vertex have answers 1 or 0, respectively. Otherwise write $e = |E|$ and $d(v) = \deg(v)$.

Make one labeled bin for each edge. At every vertex create $d(v)$ labeled items, of which one is marked. For a radix $B = 2e(q + 1) + 2$, interpret $(c_0, \ldots, c_4)$ as the integer $\sum_{k=0}^{4} c_k B^k$. Assign the following weights, where $t = 1$ denotes a marked item and $t = 0$ an unmarked one.

Object	Digits of its weight or bin total
Item at left vertex i	$(1, 0, i, 0, t)$
Item at right vertex j	$(0, 1, 0, j, 1 - t)$
Bin of edge (i, j)	$(1, 1, i, j, 1)$

The total of each coordinate over all items is less than B, so equality of weights in a bin is coordinatewise equality. The first four coordinate sums over all items agree with the bin totals by counting edge incidences. The last coordinate sums to $q + (e - q) = e$ on the item side as well.

The first two coordinates force each bin to receive one left item and one right item. The next two force their vertices to be the endpoints of the edge. The last forces their marks to agree. Thus bins containing marked items form a perfect matching. Given a perfect matching, put the marked items in the corresponding bins and permute each vertex's remaining items arbitrarily among its other incident bins. This gives

$$\mathrm{P}(a, b) = \mathrm{per}(G) \prod_{v \in L \cup R} (d(v) - 1)!.$$

The factor is computable in polynomial time. There are $2e$ items and e bins, all of weight less than B^5. The unary instance and the exact division therefore both have polynomial bit complexity. □

Proposition 4.2. *For n items of total weight N, the ordered set partition count is recoverable from n nonadaptive queries to either a nonnegative-table oracle or a binary-table oracle. Each query has unary length $O(N^2 \log(2N))$.*

Proof. Choose n distinct primes $p > N$. It suffices to sieve through the $(N + n)$th prime. The bound $p_k = O(k \log(2k))$ for the kth prime [RS62, Thm. 3] and $n \leq N$ show that we can take every $p = O(N \log(2N))$. For each prime, query the chosen oracle on $(a^{[p]}, b^{[p]})$. These lists have unary length $O(pN)$, and (2) gives $\mathrm{P}(a, b)$ modulo p from the oracle's answer.

There are at most m^n assignments. Since $m \leq N$, the product of the chosen primes strictly exceeds m^n. The Chinese remainder theorem thus recovers the unique answer in $[0, m^n]$ with

$O(n^2N^2)$ many operations from the pre-computed contingency table numbers. All arithmetic is on integers of polynomial bit length. In particular $B_p \le T_p \le (N+1)^{p^2nm}$, since the replicated table has p^2nm entries, each at most N. □

Proof of Theorem 1.1. Valiant's theorem gives #P-hardness of the 0–1 permanent [Val79]. Composing Propositions 4.1 and 4.2 gives (5) and proves hardness. For membership in #P, a verifier reads the nm entries in row order, using $\lceil \log_2(N+1) \rceil$ bits per entry for $N > 0$. For binary tables it reads one bit per cell. In either case it checks the entry bounds and all margins in polynomial time. Each table has a unique accepted encoding. Total zero and unequal totals are decided directly. □

5. Rectangular Kostka numbers

For a partition λ and a weak composition η of $|\lambda|$, let $K_{\lambda,\eta}$ count the semistandard Young tableaux of shape λ and content η. Such a tableau has η_i entries equal to i, with rows weakly increasing and columns strictly increasing. Write $\eta^{\downarrow}$ for the partition obtained by sorting the positive parts of η. Symmetry of the Schur function gives $K_{\lambda,\eta} = K_{\lambda,\eta^{\downarrow}}$ [Sta99, §7.10]. The following construction is due to Kirillov [Kir01], we repeat it here for completeness.

Theorem 5.1. *Let $a = (a_1, \dots, a_n)$ and $b = (b_1, \dots, b_m)$ be positive margins of common total N. For the rectangle $R = (N^n)$ and the content*

$$\eta = (b_1, \dots, b_m, N - a_n, \dots, N - a_1),$$

we have $\mathrm{T}(a, b) = K_{R,\eta^{\downarrow}}$.

Proof. The matrix Robinson–Schensted–Knuth correspondence [Knu70] sends each table bijectively to a pair (P, Q) of the same shape ρ, where P has content b and Q has content a. The diagram ρ has at most N columns since it has N boxes. It has at most n rows since the columns of Q are strictly increasing.

Represent a tableau by the chain of diagrams formed by its entries at most k, as k increases. Each new label adds a horizontal strip, that is, at most one box per column. Let α^k be the diagram of entries at most k in Q, and write $h_c(\alpha)$ for column c's height, using zero for an empty column. Start with the chain for P, ending at ρ. Append the diagrams $\gamma^1, \dots, \gamma^n$ defined by

$$h_c(\gamma^t) = t + h_c(\alpha^{n-t}) \qquad (1 \le c \le N, \quad 0 \le t \le n).$$

These are Young diagrams, with $\gamma^0 = \rho$ and $\gamma^n = R$. From γ^{t-1} to γ^t a column grows precisely when it did not grow from α^{n-t} to α^{n-t+1}. Thus the appended strip has size $N - a_{n-t+1}$. The full chain gives a tableau of shape R and content η.

Conversely, split such a tableau's chain after its first m labels to recover P. At the tth subsequent step each column has height at least t, since only $n - t$ steps remain to reach height n. Subtracting t from every column height of γ^t therefore recovers a Young diagram α^{n-t}. In reverse order these diagrams grow by horizontal strips of sizes $a_1, \dots, a_n$, recovering Q. This inverts the construction. Sorting the content gives the asserted equality. □

Theorem 5.2. *Computing $K_{\lambda,\mu}$ for partitions given in unary is #P-complete under deterministic polynomial-time Turing reductions, even when λ is a rectangle.*

Proof. By Theorem 5.1, an oracle for rectangular Kostka numbers evaluates each table count with a single query of the same value. Both partitions have size $nN \le N^2$, so they can be written in unary in polynomial time. Hardness follows from Theorem 1.1. For membership, guess one label per box using a fixed-width binary encoding and verify the content and the row and column inequalities. With unary input this takes polynomial time, and each tableau has exactly one accepting encoding. □

6. Littlewood–Richardson coefficients

Write $c^{\nu}_{\alpha,\lambda}$ for the coefficient of s_ν in $s_\alpha s_\lambda$. Equivalently, $c^{\nu}_{\alpha,\lambda}$ is the coefficient of s_λ in the skew Schur function $s_{\nu/\alpha}$. We spell out the Kostka-to-Littlewood–Richardson reduction of Narayanan [Nar06, Lem. 2] for an arbitrary number of rows.

Proposition 6.1. *Let λ, μ be partitions of $M > 0$ and let $r = \ell(\mu)$. Set $\nu_i = \sum_{j=i}^{r} \mu_j$ and $\alpha_i = \sum_{j=i+1}^{r} \mu_j$ for $1 \le i \le r$, omitting the final zero in α. Then*

$$K_{\lambda,\mu} = c^{\nu}_{\alpha,\lambda}.$$

Moreover $|\nu| = \sum_{j=1}^{r} j\mu_j \le rM$.

Proof. Row i of ν/α occupies columns $\alpha_i + 1, \dots, \nu_i$. Since $\alpha_i = \nu_{i+1}$ for $i < r$, distinct rows use disjoint sets of columns. A semistandard filling therefore consists of independent weakly increasing rows, giving

$$s_{\nu/\alpha} = h_{\mu_1} \cdots h_{\mu_r}.$$

By the Pieri rule, the coefficient of s_λ in this product counts chains from $\varnothing$ to λ adding successive horizontal strips of sizes $\mu_1, \dots, \mu_r$ [Sta99, Ch. 7]. Labeling the ith strip by i identifies these chains with tableaux of shape λ and content μ. Comparing coefficients proves the identity. Summing the suffix sums ν_i gives the size bound. □

Theorem 6.2. *Computing $c^{\nu}_{\alpha,\lambda}$ from three unary partitions is #P-complete under deterministic polynomial-time Turing reductions, even when λ is rectangular and ν/α has at most one box in each column.*

Proof. Apply Proposition 6.1 to the rectangular instances of Theorem 5.2. The query preserves the count, and $|\alpha| \le |\nu| \le rM \le M^2$, so all three partitions can be written in unary in polynomial time.

For membership, the Littlewood–Richardson rule counts semistandard tableaux of shape ν/α and content λ whose reading word is a lattice word [Sta99, Ch. 7]. Read each row from right to left, starting at the top. In every prefix the number of i's must be at least the number of $(i+1)$'s for every i. A verifier checks this condition, the row and column inequalities, and the content in polynomial time. With a fixed-width encoding for each label every tableau contributes one accepting computation. □

7. Schubert structure constants

Write $C^{w}_{u,v}$ for the coefficient of the Schubert polynomial $\mathfrak{S}_w$ in $\mathfrak{S}_u\mathfrak{S}_v$. These are the ordinary type A structure constants of the complete flag variety. The input consists of three permutations in one-line notation.

Theorem 7.1. *Computing $C^{w}_{u,v}$ with input the permutations w, u, v in one-line notation is #P-hard under deterministic polynomial-time Turing reductions. The restriction to triples of permutations with a unique descent at a common position k is #P-complete. Hardness holds even when one factor corresponds to a rectangular partition.*

Proof. Take an instance $c^{\nu}_{\alpha,\lambda}$ of Theorem 6.2, with $M = |\nu| = |\alpha| + |\lambda| > 0$. Set $k = \max\{\ell(\alpha), \ell(\lambda), \ell(\nu)\}$ and $r = k + M$. For each $\theta \in \{\alpha, \lambda, \nu\}$, pad θ to length k. Define $w_\theta \in S_r$ by

$$w_\theta(i) = i + \theta_{k+1-i} \quad (1 \le i \le k),$$

and list the unused elements of $[r]$ in increasing order in the remaining positions. Both blocks increase, so the only possible descent is at k. The Grassmannian identity [Mac91, (4.8)] gives

$\mathfrak{S}_{w_\theta} = s_\theta(x_1, \ldots, x_k)$. Expanding the product of the two Schur polynomials by the Littlewood–Richardson rule and comparing Schubert basis coefficients yields

$$C^{w_\nu}_{w_\alpha, w_\lambda} = c^{\nu}_{\alpha,\lambda}. \tag{6}$$

Every partition of M with at most k parts fits in the k-by-M rectangle, so this comparison uses permutations in S_r. Since $r \leq 2M$, writing the three permutations uses $O(M \log(2M))$ bits and takes polynomial time in the unary input. This gives the claimed count-preserving reduction, including the rectangular restriction. Conversely the first k entries of each permutation recover its partition by the same formula. The partitions have at most $k(r-k)$ boxes each, so the Littlewood–Richardson verifier proves #P membership on the restricted family. The identity permutation represents the empty partition. □

Pak [Pak24, §13.4(10)], Pak and Robichaux [PR25b, Conj. 1.2] explicitly formulate the hardness conjecture. Mukhopadhyay and Qiao [MQ17, p. 885] and Adve, Robichaux, and Yong [ARY19, §4] inferred hardness from Narayanan's binary-input theorem. Ikenmeyer, Pak, and Panova [IPP24, §5.2] identified the encoding issue that binary encoding is not the right setting for Schuberts and recorded the implication from unary hardness. For example, the one-row partition (L) with $L = 2^t$ takes $t+1$ bits to specify its part, while its Grassmannian permutation at $k = 1$ requires at least $L+1$ entries. The unary bound above supplies the missing step.

Mukhopadhyay and Qiao also give an output-sensitive interpolation algorithm with a #P oracle [MQ17]. Adve, Robichaux, and Yong prove #P-completeness for the monomial coefficient $[x^\beta]\mathfrak{S}_w$ when w is specified by its code [ARY21]. For general triples, Morales's argument [Pak24, Prop. 10.2] and the signed formulas of Pak and Robichaux [PR25a, PR26] place exact evaluation in GapP. Together with Theorem 7.1, this gives polynomial-time Turing equivalence with the 0–1 permanent. General membership in the discrete class #P remains open, in connection with Stanley's request for a positive combinatorial rule [Sta00, Problem 11]. Positive combinatorial formulas are known only for special classes of triples of permutations, see [BKPT16, KZJ26c, KZJ26a, KZJ26b, Hua23, Sot96, MPP14, Yu26, GZ26].

8. AI disclosure

The human author is one of the authors posing the original question of the complexity of computing the number of contingency tables, Kostka, Littlewood-Richardson with inputs in unary. Besides a fundamental question it theoretical computer science, this would also answer the question whether computing Schubert coefficients with inputs given by the 3 permutations is #P-hard. Following the recent AI breakthroughs the author decided to check this problem with chatGPT, which answered it very quickly. Prompting the model to find it answered previously anywhere on the web returned negative results. Since this solution was so short and elegant, and the problem important in both TCS and algebraic combinatorics, the human author decided that it needs to be made known to the community. However, we believe that paper production by pushing buttons and taking credit for it is pointless and destructive to our peer review system, to say the least. Research paper production is suffering from hyperinflation and is effectively devalued. Therefore we decided that such work should not be posted on the arXiv, but instead go to the hexagonmath repository. Nevertheless, the human author has written some and edited other parts of the text, and verified everything carefully.

References

[ARY19] Anshul Adve, Colleen Robichaux, and Alexander Yong, *Vanishing of Littlewood–Richardson polynomials is in* P, Computational Complexity **28** (2019), no. 2, 241–257.

[ARY21] ———, *An efficient algorithm for deciding vanishing of Schubert polynomial coefficients*, Advances in Mathematics **383** (2021), 107669.

[BKPT16] Anders Skovsted Buch, Andrew Kresch, Kevin Purbhoo, and Harry Tamvakis, *The puzzle conjecture for the cohomology of two-step flag manifolds*, Journal of Algebraic Combinatorics **44** (2016), no. 4, 973–1007.

[DKM97] Martin Dyer, Ravi Kannan, and John Mount, *Sampling contingency tables*, Random Structures & Algorithms **10** (1997), no. 4, 487–506.

[GZ26] Yibo Gao and Hai Zhu, *Boolean Schubert structure coefficients*, Journal of Algebra **688** (2026), 344–362.

[Hua23] Daoji Huang, *Schubert products for permutations with separated descents*, International Mathematics Research Notices **2023** (2023), no. 20, 17461–17493.

[IPP24] Christian Ikenmeyer, Igor Pak, and Greta Panova, *Positivity of the symmetric group characters is as hard as the polynomial time hierarchy*, International Mathematics Research Notices **2024** (2024), no. 10, 8442–8458.

[Kir01] Anatol N. Kirillov, *Ubiquity of Kostka polynomials*, Physics and Combinatorics 1999, World Scientific, Singapore, 2001, pp. 85–200.

[Knu70] Donald E. Knuth, *Permutations, matrices, and generalized Young tableaux*, Pacific Journal of Mathematics **34** (1970), no. 3, 709–727.

[KZJ26a] Allen Knutson and Paul Zinn-Justin, *Schubert puzzles and integrability I: invariant trilinear forms*, Communications of the American Mathematical Society **6** (2026), 1–67.

[KZJ26b] ———, *Schubert puzzles and integrability II: multiplying motivic Segre classes*, Communications of the American Mathematical Society **6** (2026), 68–152.

[KZJ26c] ———, *Schubert puzzles and integrability III: separated descents*, The Electronic Journal of Combinatorics **33** (2026), no. 4, Paper P4.39.

[Mac91] I. G. Macdonald, *Notes on Schubert polynomials*, Publications du LaCIM, vol. 6, Université du Québec à Montréal, 1991.

[MH13] Jeffrey W. Miller and Matthew T. Harrison, *Exact sampling and counting for fixed-margin matrices*, The Annals of Statistics **41** (2013), no. 3, 1569–1592.

[MPP14] Karola Mészáros, Greta Panova, and Alexander Postnikov, *Schur times Schubert via the Fomin–Kirillov algebra*, The Electronic Journal of Combinatorics **21** (2014), no. 1, P1.39.

[MQ17] Priyanka Mukhopadhyay and Youming Qiao, *Sparse multivariate polynomial interpolation on the basis of Schubert polynomials*, Computational Complexity **26** (2017), no. 4, 881–909.

[Nar06] Hariharan Narayanan, *On the complexity of computing Kostka numbers and Littlewood–Richardson coefficients*, Journal of Algebraic Combinatorics **24** (2006), no. 3, 347–354.

[Ope26a] OpenAI, *Exact uniform sampling of contingency tables with arbitrary margins*, OpenAI Math Release preprint, September 2026.

[Ope26b] ———, *An FPRAS for cell-bounded contingency tables*, OpenAI Math Release preprint, September 2026.

[Pak24] Igor Pak, *What is a combinatorial interpretation?*, Open Problems in Algebraic Combinatorics (Christine Berkesch, Benjamin Brubaker, Gregg Musiker, Pavlo Pylyavskyy, and Victor Reiner, eds.), Proceedings of Symposia in Pure Mathematics, vol. 110, American Mathematical Society, Providence, RI, 2024, pp. 191–260.

[Pan25] Greta Panova, *Computational complexity in algebraic combinatorics*, Current Developments in Mathematics **2023-2024** (2025), no. 1, 241–280.

[PR25a] Igor Pak and Colleen Robichaux, *Signed combinatorial interpretations in algebraic combinatorics*, Algebraic Combinatorics **8** (2025), no. 2, 495–519.

[PR25b] ———, *Vanishing of Schubert coefficients*, arXiv:2412.02064v3, 2025.

[PR26] ———, *Signed puzzles for Schubert coefficients*, Algebraic Combinatorics **9** (2026), 1479–1495.

[RS62] J. Barkley Rosser and Lowell Schoenfeld, *Approximate formulas for some functions of prime numbers*, Illinois Journal of Mathematics **6** (1962), no. 1, 64–94.

[Sot96] Frank Sottile, *Pieri's formula for flag manifolds and Schubert polynomials*, Annales de l'Institut Fourier **46** (1996), no. 1, 89–110.

[Sta99] Richard P. Stanley, *Enumerative combinatorics. vol. 2*, Cambridge Studies in Advanced Mathematics, vol. 62, Cambridge University Press, Cambridge, 1999.

[Sta00] ———, *Positivity problems and conjectures in algebraic combinatorics*, Mathematics: Frontiers and Perspectives, American Mathematical Society, Providence, RI, 2000, pp. 295–319.

[Val79] L. G. Valiant, *The complexity of computing the permanent*, Theoretical Computer Science **8** (1979), no. 2, 189–201.

[Yu26] Tianyi Yu, *Two-row Schur times Schubert via parallel classes of positroid varieties*, arXiv:2611.07351, 2026.

University of Southern California, Los Angeles, CA