COMBINATORIAL INEQUALITIES

IGOR PAK⋆

Combinatorics has always been a battleground of tools and ideas. That's why it's so hard to do, or even define. The inequalities are a particularly interesting case study as they seem to be both the most challenging and the least explored in Enumerative and Algebraic Combinatorics. Here are a few of my favorites, with some backstories.[1]

We start with *unimodality of binomial coefficients:*

$$\binom{n}{k-1} \le \binom{n}{k}, \quad \text{for all } 1 \le k \le n/2. \tag{1}$$

This is both elementary and well known – the proof is an easy calculation. But ask yourself the following natural question: does the difference $B(n,k) := \binom{n}{k} - \binom{n}{k-1}$ count anything interesting? It should, of course, right? Imagine there is a natural injection

$$\psi : \binom{[n]}{k-1} \to \binom{[n]}{k}$$

from $(k-1)$-subsets to k-subsets of $[n]$, where $[n] := \{1,\ldots,n\}$. Then $B(n,k)$ can be described as the number of k-subsets of $[n]$ that are not in the image of ψ, as good answer as any. But how do you construct the injection ψ?

Let us sketch the construction based on the classical *reflection principle* for the *ballot problem*, which goes back to the works of Bertrand and André in 1887. Start with a $(k-1)$-subset X of $[n]$, and let ℓ be the smallest integer s.t. $\bigl|X \cap [2\ell+1]\bigr| = \ell$. Such ℓ exists since $k \le n/2$. Define

$$\psi(X) := \bigl(X \smallsetminus [2\ell+1]\bigr) \cup \bigl([2\ell+1] \smallsetminus X\bigr).$$

Observe that $|\psi(X)| = k$ and check that ψ is the desired injection. This gives an answer to the original question: $B(n,k)$ is the number of k-subsets $Y \subset [n]$, s.t. $\bigl|Y \cap [m]\bigr| \le m/2$ for all m.

At this point you might be in disbelief in me dwelling on the easy inequality (1). Well, it only gets harder from here. Consider, e.g., the following question: Does there exist an injection ψ as above, s.t. $X \subset \psi(X)$ for all $X \in \binom{[n]}{k-1}$? We leave it to the reader as a challenge.[2]

There is also a curious connection to Algebraic Combinatorics: $B(n,k) = f^{(n-k,k)}$, the dimension of the irreducible S_n-module corresponding to the partition $(n-k,k)$. To understand how this could happen, think of both sides of (1) as dimensions of permutation representations of S_n. Turn both sides into vector spaces and modify ψ accordingly, to make it an S_n-invariant linear map. This would make it more natural and uniquely determined. As a consequence, we obtain a combinatorial interpretation $B(n,k) = \bigl|\mathrm{SYT}(n-k,k)\bigr|$, the number of standard Young tableaux of shape $(n-k,k)$, a happy outcome in every way.

Consider now *unimodality of Gaussian coefficients:*

$$p(n,k,\ell-1) \le p(n,k,\ell), \quad \text{for all } 1 \le \ell \le k(n-k)/2, \text{ where} \tag{2}$$

⋆Department of Mathematics, UCLA, Los Angeles, CA 90095, USA. Email: pak@math.ucla.edu.

[1] To save space and streamline the exposition we omit all of the many references. The interested reader can find them in the expanded version of this article at the author's website: https://tinyurl.com/y26efyoj.

[2] We give the answer in the expanded version, see above.

$p(n,k,\ell)$ is the number of integer partitions $\lambda \vdash \ell$ that fit into a $k \times (n-k)$ rectangle, i.e. λ has parts of size at most $(n-k)$, and has at most k parts. To understand the context of this inequality, recall:

$$\sum_{\ell=0}^{k(n-k)} p(n,k,\ell)\, q^\ell \, = \, \binom{n}{k}_q \, := \, \frac{(n!)_q}{(k!)_q \cdot \big((n-k)!\big)_q}\,, \quad \text{where} \quad (n!)_q \, := \, \prod_{i=1}^{n} \frac{q^i-1}{q-1}\,.$$

To connect this to (1), note that $\binom{n}{k}_1 = \binom{n}{k}$, and that $\binom{n}{k}_q$ is the number of k-subspaces of $\mathbb{F}_q^n$. In (2), we view $\binom{n}{k}_q$ as a polynomial in q and compare its coefficients. Now, the Schubert cell decomposition of the Grassmannian over $\mathbb{F}_q$, or a simple induction can be used to give the partition interpretation.

The inequality (2) is no longer easy to prove. Conjectured by Cayley in 1856, it was established by Sylvester in 1878; the original paper is worth reading even if just to see how pleased Sylvester was with his proof. In modern language, Sylvester defined the $\mathfrak{sl}_2(\mathbb{C})$ action on certain homogeneous polynomials and the result follows from the highest weight theory (in its simplest form for $\mathfrak{sl}_2$).

Let's continue with the questions as we did above. Consider the difference $C(n,k,\ell) := p(n,k,\ell) - p(n,k,\ell-1)$. Does $C(n,k,\ell)$ count anything interesting? Following the pattern above, wouldn't it be natural to define some kind of nice injection from partitions of size $(\ell-1)$ to partition of size ℓ, by simply adding a corner square according to some rule? That would be an explicit combinatorial (as opposed to algebraic) version of Sylvester's approach.

Unfortunately we don't know how to construct such a nice injection. It's just the first of the many frustrations one encounters with algebraic proofs. Most of them are simply too rigid to be "combinatorialized". It doesn't mean that there is no combinatorial interpretation for $C(n,k,\ell)$ at all. There is one very uninteresting interpretation due to Panova and myself, based on a very interesting (but cumbersome) identity by O'Hara. Also, from the Computer Science point of view, it is easy to show that $C(n,k,\ell)$ as a function is in $\#\mathsf{P}$. We leave it to the reader to figure out why (or what does that even mean).

To finish this story, we should mention Stanley's 1989 approach to (2) using finite group actions. More recently, Panova and I introduced a different technique based on properties of the *Kronecker coefficients* of S_n, via the equality $C(n,k,\ell) = g\big((n-k)^k, (n-k)^k, (n-\ell,\ell)\big)$. Here the Kronecker coefficients $g(\lambda,\mu,\nu)$ can be defined as structure constants for products of S_n characters: $\chi^\mu \chi^\nu = \sum_\lambda g(\lambda,\mu,\nu)\,\chi^\lambda$. Both approaches imply stronger inequalities than (2), but neither gets us closer to a simple injective proof.

We turn now to *log-concavity of independent sets:*

$$a_{k-1}(M) \cdot a_{k+1}(M) \, \le \, a_k(M)^2, \quad \text{where} \tag{3}$$

$a_k(M)$ is the number of independent k-subsets of a matroid M. Note that the log-concavity implies unimodality, and in the special case of a *free matroid* (all elements are independent) this gives (1).

The inequality (3) is a celebrated recent result by Adiprasito, Huh and Katz (2018), which showed that a certain "cohomology ring" associated with M satisfies the hard Lefschetz theorem and the Hodge–Riemann relations. This resolved conjectures by Welsh and Mason (1970s).

It would be naïve for us to ask for a direct combinatorial proof via an injection, or by some other elementary means. For example, Stanley in 1981 used the Aleksandrov–Fenchel inequalities in convex geometry to prove that the log-concavity is preserved under taking truncated sum with a free matroid, already an interesting but difficult special case proved by inherently non-combinatorial means.

There is also a Computational Complexity version of the problem which might be of interest. Let $A(k,M) := a_k(M)^2 - a_{k-1}(M) \cdot a_{k+1}(M)$. Does $A(k,M)$ count any set of combinatorial objects?

For the sake of clarity, let $G=(V,E)$ be a simple connected graph and M the corresponding matroid, i.e. bases in M are spanning trees in G. Then $a_k(M)$ is the number of spanning forests in G with k edges. Note that computing $a_k(M)$ is #P-complete in full generality. Therefore, computing $A(k,M)$ is #P-hard.

Now, $A(k,M)$ is in GapP, i.e. equal to the difference of two #P-functions. Does $A(k,M)$ lie in #P? This seems unlikely, but the current state of art of Computational Complexity doesn't seem to provide us with tools to even approach a negative solution.

To fully appreciate the last example, consider the *log-concavity of matching numbers:*

$$m_{k-1}(G) \cdot m_{k+1}(G) \le m_k(G)^2, \quad \text{where} \tag{4}$$

$m_k(G)$ is the number of k-matchings in a simple graph $G=(V,E)$, i.e. k-subsets of edges which are pairwise disjoint. For example, $m_n(K_{2n}) = (2n-1)\cdots 3\cdot 1$. While perfect matchings don't necessarily define a matroid, they do have a similar flavor from a Combinatorial Optimization point of view. The inequality (4) goes back to Heilmann and Lieb (1972) and is a rare case when the injection strategy works well. The following argument is due to Krattenthaler (1996).

Take a $(k-1)$-matching β whose edges we color *blue* and a $(k+1)$-matching γ whose edges we color *green*. The union $\beta\cup\gamma$ of these two sets of edges splits into connected components, which are either paths or cycles, all alternately colored. Ignore for the time being all cycles and paths of even lengths. Denote by $(r-1)$ the number of odd-length paths which have extra color blue. There are then $(r+1)$ odd-length paths which have extra color green.

Now, allow switching colors in any of the $2r$ odd-length paths. After recoloring, we want to have r odd-length paths extra color blue and the same with green. This amounts to a constructive injection from $(r-1)$-subsets of $[2r]$ to r-subsets of $[2r]$, which we already know how to do as a special case of proving (1).

We leave to the reader the problem of finding an explicit combinatorial interpretation for $M(k,G) := m_k(G)^2 - m_{k-1}(G)\cdot m_{k+1}(G)$, proving that this function is in #P. Note that computing $m_k(G)$ is famously #P-complete, which implies that so is $M(k,G)$. This makes the whole connection to Computational Complexity even more confusing. What exactly makes matchings special enough for this argument to work?

If there is any pattern to the previous examples, it can be summarized as follows: the deeper one goes in an algebraic direction, the more involved are the inequalities and the less of a chance of a combinatorial proof. To underscore this point, consider the following three *Young tableaux inequalities*:

$$\big(f^\lambda\big)^2 \le n!, \qquad \big(c^\lambda_{\mu\nu}\big)^2 \le \binom{n}{k}, \qquad c^\lambda_{\mu\nu} \le c^\lambda_{\mu\vee\nu,\,\mu\wedge\nu}\,, \qquad \text{for all } \lambda\vdash n,\ \mu\vdash k,\ \nu\vdash n-k. \tag{5}$$

Here $f^\lambda = \big|\mathrm{SYT}(\lambda)\big|$ is the number of standard Young tableaux of shape λ, equal to the dimension of the corresponding irreducible S_n-module as above. Similarly, $c^\lambda_{\mu\nu} = \big|\mathrm{LR}(\lambda/\mu,\nu)\big|$ is the *Littlewood–Richardson coefficient*, equal to the number of Littlewood–Richardson tableaux of shape λ/μ and weight ν. It can be defined as a structure constant for products of Schur functions: $s_\mu\, s_\nu = \sum_\lambda c^\lambda_{\mu\nu}\, s_\lambda$. Finally, $\mu\vee\nu$ and $\mu\wedge\nu$ denote the union and intersection, respectively, of the corresponding Young diagrams.

Now, the first inequality in (5) is trivial algebraically, but its combinatorial proof is highly nontrivial – it is a restriction of the RSK correspondence. The second inequality is quite recent and follows easily from the definition and the Frobenius reciprocity. We believe it is unlikely that there is a combinatorial injection, even though there is a nice double counting argument.

Finally, the third inequality in (5) is a corollary of the powerful inequality by Lam, Postnikov and Pylyavskyy (2007) using the curious Temperley–Lieb immanant machinery. The key ingredient in the proof is Haiman's theorem which in turn uses the Kazhdan–Lusztig conjecture proven by Beilinson–Bernstein and Brylinski–Kashiwara. While stranger things have happened, we would be very surprised if this inequality had a simple combinatorial proof.

We conclude on a positive note, with a combinatorial inequality where everything works as well as it possibly could. Consider the following *majorization property of contingency tables*:

$$\mathrm{T}(\mathbf{a},\mathbf{b}) \le \mathrm{T}(\mathbf{a}',\mathbf{b}') \quad \text{for all } \mathbf{a}'\trianglelefteq\mathbf{a},\ \mathbf{b}'\trianglelefteq\mathbf{b}. \tag{6}$$

Here $\mathbf{a} = (a_1, \ldots, a_m)$, $a_1 \geq \ldots \geq a_m > 0$, and $\mathbf{b} = (b_1, \ldots, b_n)$, $b_1 \geq \ldots \geq b_n > 0$, are two integer sequences with equal sum:

$$\sum_{i=1}^{m} a_i \,=\, \sum_{j=1}^{n} b_j \,=\, N.$$

A *contingency table* with margins $(\mathbf{a}, \mathbf{b})$ is an $m \times n$ matrix of non-negative integers whose i-th row sums to a_i and whose j-th column sums to b_j, for all $i \in [m]$ and $j \in [n]$. $\mathrm{T}(\mathbf{a}, \mathbf{b})$ denotes the number of all such matrices. Finally, for sequences $\mathbf{a}$ and $\mathbf{a}'$ with the same sum, we write $\mathbf{a} \trianglelefteq \mathbf{a}'$ if $a_1 \leq a_1'$, $a_1 + a_2 \leq a_1' + a_2'$, $a_1 + a_2 + a_3 \leq a_1' + a_2' + a_3'$, $\ldots$ In other words, the inequality (6) says that there are more contingency tables when the margins are more evenly distributed.

Contingency tables can be viewed as adjacency matrices of bipartite multi-graphs with given degree distribution. They play an important role in Statistics and Network Theory. We learned the inequality (6) from a paper by Barvinok (2007), but it feels like something that should have been known for decades.

Now, I know two fundamentally different proofs of (6). The first is an algebraic proof using Schur functions which amounts to proving the following standard inequality for Kostka numbers: $K_{\lambda\mu} \leq K_{\lambda\nu}$ for all $\mu \trianglerighteq \nu$, where $K_{\lambda\mu}$ is the number of semistandard Young tableaux of shape λ and weight μ. This inequality can also be proved directly, so combined with the RSK we obtain an injective proof of (6).

Alternatively, one can prove the inequality directly for $2 \times n$ rectangles and $(+1, -1)$ changes in row (column) sums. Combining these injections together gives a cumbersome, yet explicit injection. In principle, either of the two approaches can then be used to give a combinatorial interpretation for $\mathrm{T}(\mathbf{a}', \mathbf{b}') - \mathrm{T}(\mathbf{a}, \mathbf{b})$.

In conclusion, let us note that we came full circle. Let $m = 2$, $a_1 = n - k + 1$, $a_2 = k - 1$, $a_1' = n - k$, $a_2' = k$, and $b_1 = \ldots = b_n = b_1' = \ldots = b_n' = 1$. Observe that $\mathrm{T}(\mathbf{a}, \mathbf{b}) = \binom{n}{k-1}$ and $\mathrm{T}(\mathbf{a}', \mathbf{b}') = \binom{n}{k}$. The inequality (1) is a special case of (6) then.