Cycles in Mallows random permutations

Tobias Müller

University of Groningen

We study random permutations of $1, \ldots, n$ drawn at random according to the Mallows distribution. For $n \in \mathbb{N}$ and $q > 0$, the distribution Mallows(n, q) samples a random permutation Π_n of $1, \ldots, n$ in such a way that each has probability proportional to $q^{\mathrm{inv}(\pi)}$, where $\mathrm{inv}(\pi)$ is the number of *inversions*. That is, pairs $1 \leq i < j \leq n$ for which $\pi(i) > \pi(j)$. In a formula:

$$\mathbb{P}(\Pi_n = \pi) = \frac{q^{\mathrm{inv}(\pi)}}{\sum_{\sigma \in S_n} q^{\mathrm{inv}(\sigma)}}, \tag{0.1}$$

for all $\pi \in S_n$ where S_n denotes the set of permutations of $1, \ldots, n$.

This distribution was introduced in the late fifties by C.L. Mallows in the context of "statistical ranking models" and has since been studied in connection with a diverse range of topics.

In the present work we will consider the cycle counts. That is, for ℓ fixed we study the vector $(C_1(\Pi_n), \ldots, C_\ell(\Pi_n))$ where $C_i(\pi)$ denotes the number of cycles of length i in π and Π_n is sampled according to the Mallows distribution.

When $q = 1$ then the Mallows distribution is simply the uniform distribution on S_n. A classical result going back to Kolchin and Goncharoff states that in this case, the vector of cycle counts tends in distribution to a vector of independent Poisson random variables, with means $1, \frac{1}{2}, \frac{1}{3}, \ldots, \frac{1}{\ell}$.

Surprisingly, the problem of finding analogues of this result for $q \neq 1$ has largely escaped attention until now. In the talk, I plan to discuss our proof of the fact that if $0 < q < 1$ is fixed and $n \to \infty$ then the cycle counts have linear means and the vector of cycle counts can be suitably rescaled to tend to a joint Gaussian distribution. Our results also show that when $q > 1$ there is a striking difference between the behaviour of the even and the odd cycles. The even cycle counts still have linear means and when properly rescaled tend to a multivariate Gaussian distribution, while for the odd cycle counts on the other hand, the limiting behaviour depends on the parity of n when $q > 1$.

Time permitting, I may also discuss some results on the (probabilities of) properties of permutations that can be expressed in first order logic.

(Based on joint work with Jimmy He, Fiona Skerman and Teun Verstraaten)