Probabilistic preference learning with the Mallows rank model
Valeria Vitelli, Øystein Sørensen, Marta Crispino, Arnoldo Frigessi and Elja Arjas

Supplementary Material

1. Tuning the proposal distributions parameters

We study the effect of the leap-and-shift proposal on the acceptance ratio for $\boldsymbol{\rho}$, and the tuning of its parameter L in relation to MCMC convergence. Also the role of parameter σ_α in the log-normal proposal for α, is briefly explored. We use simulated data.

Data were generated from the Mallows model with footrule distance, $\alpha_{\text{true}} = 2$ and $\boldsymbol{\rho}_{\text{true}} = (1, \dots, n)$. Two scenarios were used, with $n = 20$ and $n = 50$, because the choice of L would likely depend on the number of items. For generating the data, we run our MCMC sampler (see Appendix C) for 10^5 burn-in iterations, and collected one sample every 100 iterations after that. We collected samples from $N = 500$ assessors. The data analyses were carried out by using the same distance as in the data generation (footrule), and the MCMC was run for 10^6 iterations after 10^5 iterations of burn-in, with a 1 to 100 thinning for α. 10 different chains were started from random points of the parameter space, and posterior inference was based on merging the results from these chains, as the MCMC converged to the same limit. The same analyses were also done for the Kendall distance, on data generated by using the `PerMallows` R package (Irurozki et al., 2016). Equivalent results (not shown) as for the footrule were obtained.

In MCMC, we needed to control for (i) mixing, aiming at an acceptance rate of approximately 1/3 for each parameter (Gelman et al., 1996; Roberts et al., 1997), and (ii) autocorrelation, monitoring the Integrated Autocorrelation Time (IAT) τ (Green and Han, 1992). Since $\boldsymbol{\rho}$ is multivariate, we monitored the IAT for each component of $\boldsymbol{\rho}$.

As expected, the acceptance rate η_ρ of proposals for $\boldsymbol{\rho}$ decreases with increasing L, and η_ρ depends also on the value of n (Figure 1, top panels). Based on the results shown in Figure 1 (bottom panels), we propose as a rule of thumb that L should be set equal to $n/5$. This choice seems reasonable also from the perspective of η_ρ (Figure 1, upper panels). The acceptance rate η_α of proposals for α decreases with increasing σ_α (Table 1). Aiming at a value of η_α close to 1/3 sets us also close to the minimal value of τ_α. In the case of $n = 20$ values close to 0.2 appear to be good choices for σ_α, while for $n = 50$ values near 0.1 might be slightly preferred.

2. Asymptotic behavior of $Z_n(\alpha)$ for $n \to \infty$, Mukherjee (2016)

In this Section we summarize the proposal in Mukherjee (2016) for computing the limit of $Z_n(\alpha)$ for $n \to \infty$, which we have previously denoted by $Z_{\lim}(\alpha)$. We need some additional notation. Consider the Mallows model (1), with $\boldsymbol{\rho} = \mathbf{1}_n$. If we choose $f(x, y) = -|x - y|$ then $\sum_{i=1}^n f(i/n, R_i/n) = -\frac{1}{n} \sum_{i=1}^n |i - R_i|$, which is the exponential in a Mallows footrule model term, before multiplying with α. Choosing $f(x, y) = -(x - y)^2$ leads to a Mallows Spearman

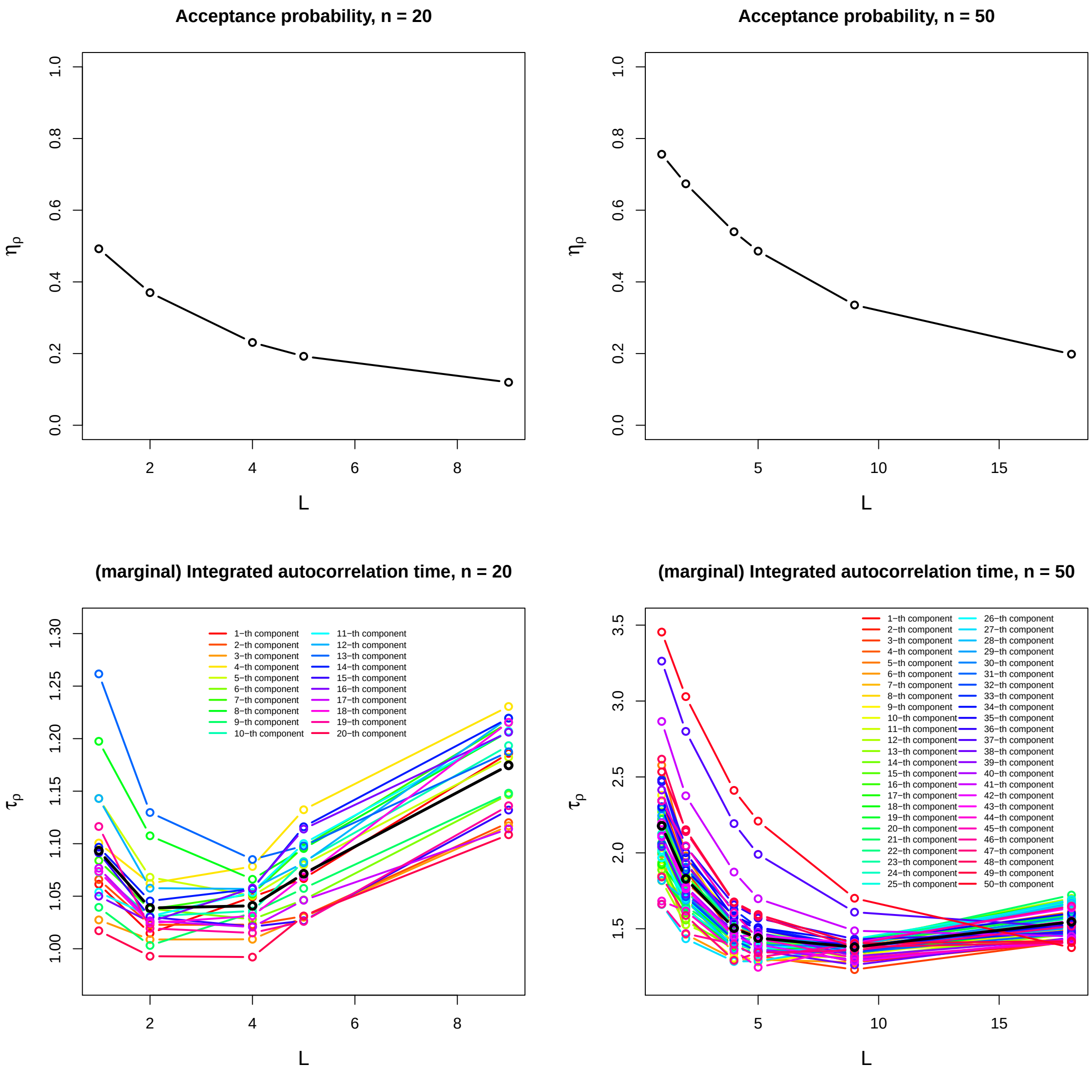

Figure 1: Results of the simulations described in Section 1. Top panels: acceptance probability η_ρ of $\boldsymbol{\rho}$ along MCMC iterations; bottom panels: marginal IAT τ_ρ of $\boldsymbol{\rho}$. Left and right panels show the results when $n = 20$ and 50, respectively.

model. Here we denote by $Z_n(f, \alpha)$ the partition function to make the dependence on the chosen distance f explicit, and its limit for $n \to \infty$ is $Z_{\lim}(f, \alpha)$. Finally, $\mathcal{M}(n, f, \alpha)$ is the Mallows model of the form above.

The main result in Mukherjee (2016) gives the limit of $Z_n(f, \alpha)$ as $n \to \infty$. It is stated as follows:

	$n = 20$		$n = 50$	
σ_α	η_α	τ_α	η_α	τ_α
0.01	0.93	4.32	0.88	3.83
0.02	0.86	3.6	0.76	3.4
0.05	0.67	2.67	0.51	2.64
0.1	0.47	2.65	**0.3**	**2.41**
0.2	**0.27**	**2.24**	0.16	2.2
0.5	0.11	2.53	0.07	2.74

Table 1: Results of the simulations described in Section 1. Acceptance probability η_α and IAT τ_α of α along MCMC iterations, for two simulations with $n = 20$ and 50. In each row, the value of σ_α (standard deviation of the log-normal proposal for α) used in the MCMC.

Theorem 1 *(Mukherjee, 2016, Theorem 1.5). For any continuous function f consider the probability model $\mathcal{M}(n, f, \alpha)$, and let $\alpha \in \mathbb{R}$ be fixed. Then*

$$\lim_{n\to\infty} \frac{Z_n(f,\alpha) - Z_n(0)}{n} = Z_{\lim}(f,\alpha) := \sup_{\mu\in\mathcal{Q}} \{\alpha\mu[f] - D(\mu||u)\},$$

where u is the uniform distribution on the unit square, $\mu[f] := \int f d\mu$ is the expectation of f with respect to the measure μ, $D(\cdot||\cdot)$ is the Kullback-Leibler divergence, $\mathcal{Q}$ is the space of all probability distributions on the unit square with uniform marginals, and $Z_n(0) = Z_n(f, 0) =$ $\log(n!)$.

As optimization over the infinite dimensional space $\mathcal{Q}$ can be hard, a second result provided in Mukherjee (2016) gives an iterative algorithm which can be used to compute a numerical approximation of $Z_{\lim}(f, \alpha)$.

Theorem 2 *IPFP - Iterative Proportional Fitting Procedure (Mukherjee, 2016, Theorem 1.9). Define a sequence of $k \times k$ matrices by setting $B_0(r, s) := e^{\alpha f(r/k,s/k)}$ for $1 \leq r, s \leq k$, and*

$$B_{2m+1}(r,s) := \frac{B_{2m}(r,s)}{k\sum_{l=1}^{m} B_{2m}(r,l)}, \quad B_{2m+2}(r,s) := \frac{B_{2m+1}(r,s)}{k\sum_{l=1}^{m} B_{2m+1}(l,s)}.$$

Then, there exists a matrix $A_{k,\alpha} \in \mathcal{Q}_k$ such that $\lim_{m\to\infty} B_m = A_{k,\alpha}$. This implies

$$Z_{\lim}(f,\alpha) = \lim_{k\to\infty} \lim_{m\to\infty} \left\{ \alpha \sum_{i,j=1}^{k} f(i/k, j/k) B_m(i,j) - 2\log k - \sum_{i,j=1}^{k} B_m(i,j) \log B_m(i,j) \right\}.$$

Due to the construction of the results in Mukherjee (2016), the role of the permutation dimension in the limit is played by k, the dimension of the grid approximating the continuous domain where the limit is computed. Hence, it is sufficient to fix k large enough in order for the continuous approximation to be reasonable (we fix $k = 10^3$, following Mukherjee (2016, Section 2)).

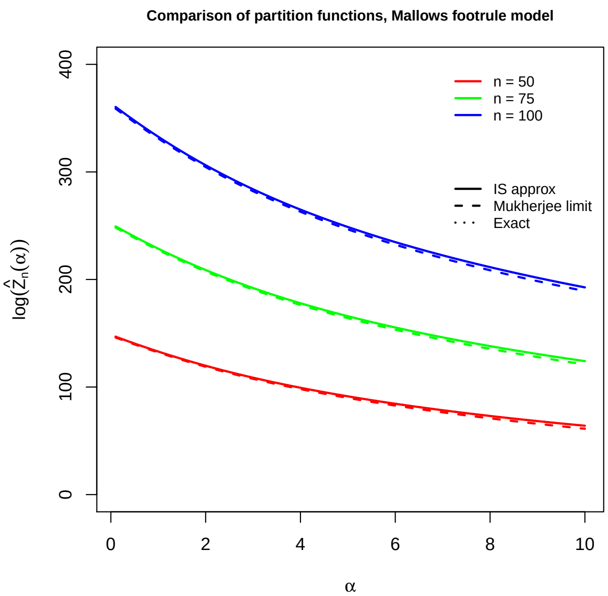

Figure 2: A comparison of different approaches to compute the partition function $Z_n(\alpha)$ for the Mallows footrule model. Different colors refer to different values of n, and different line types to different strategies for computing $Z_n(\alpha)$, as stated in the legend (note that only the IS approximation and the Mukherjee limit are available for $n = 100$).

With the aim of checking whether the limit $Z_{\text{lim}}(f, \alpha)$ is a good approximation for $Z_n(f, \alpha)$ for reasonably large values of n, we used the IPFP with $m = 10^4$ iterations (after verifying in different situations that the IPFP had typically already converged after 10^3 iterations; not shown). Finally, as it is evident from the form of the limit in Theorem 1, once we compute $Z_{\text{lim}}(f, \alpha)$ we have to rescale it in the following way

$$Z_{\text{lim},n}(f, \alpha) = n \cdot Z_{\text{lim}}(f, \alpha) + Z_n(0). \tag{1}$$

A comparison of $Z_{\text{lim},n}(f, \alpha)$ rescaled as in (1) to the IS approximation $\hat{Z}_n(f, \alpha)$ for $n = 50, 75, 100$ is shown in Figure 2. For these cases, the asymptotic approximation is quite close to the true values of the partition function. This is very useful in applications where n is so large that even the importance sampling approximation is computationally unfeasible, and thus using the limiting partition function $Z_{\text{lim}}(f, \alpha)$ for approximating $Z_n(f, \alpha)$ is an excellent alternative.

3. Additional Figures from Sections 3.3 and 4.4

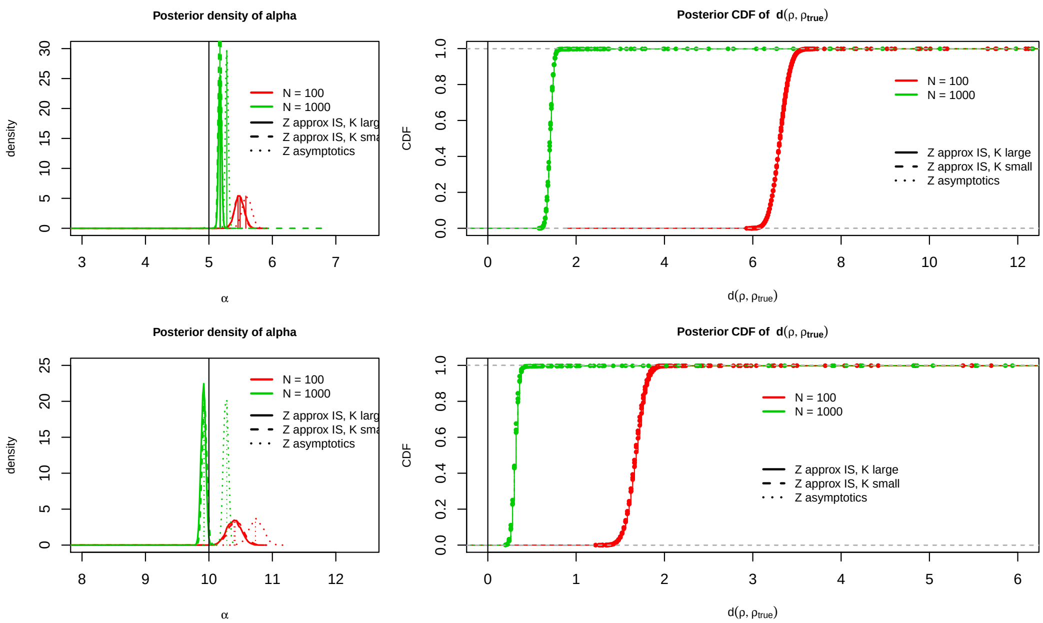

Figure 3: Results of the simulations described in Section 3.3, when $n = 100$. Left, posterior density of α (the black vertical line indicates α_{true}) obtained for various choices of N (different colors), and when using different approximations to the partition function (different line types), as stated in the legend. Right, posterior CDF of $d(\boldsymbol{\rho}, \boldsymbol{\rho}_{\text{true}})$ in the same settings. First row: $\alpha_{\text{true}} = 5$; Second row: $\alpha_{\text{true}} = 10$.

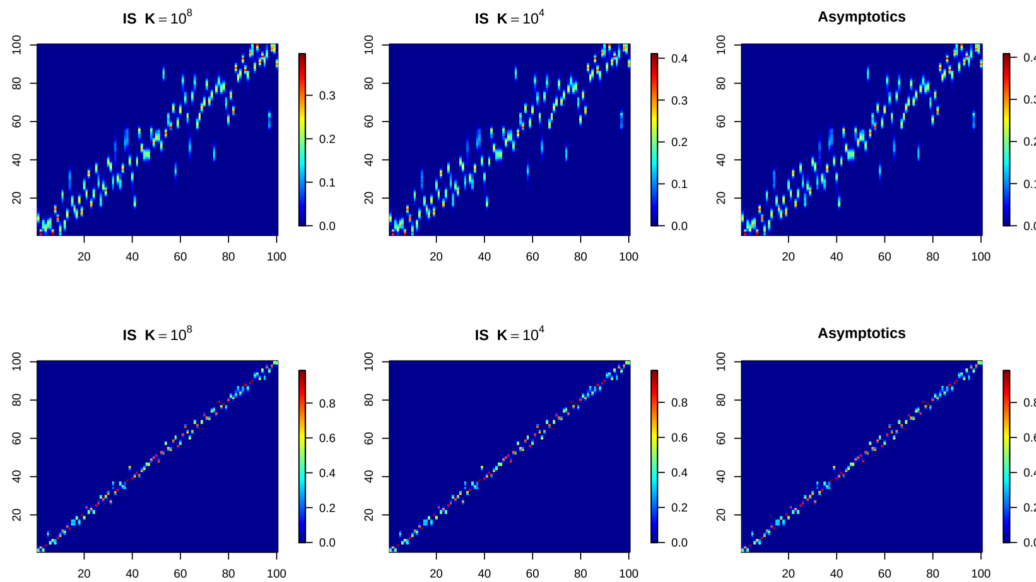

Figure 4: Results of the simulations described in Section 3.3, when $n = 100$ and $\alpha_{\text{true}} = 5$. In each heatplot, posterior marginal distribution of $\boldsymbol{\rho}$. From left to right, results obtained with the IS approximation $\hat{Z}_n^K(\alpha)$ with $K = 10^8$, with the IS approximation $\hat{Z}_n^K(\alpha)$ with $K = 10^4$, and with $Z_{\text{lim}}(\alpha)$. First row: $N = 100$; Second row: $N = 1000$.

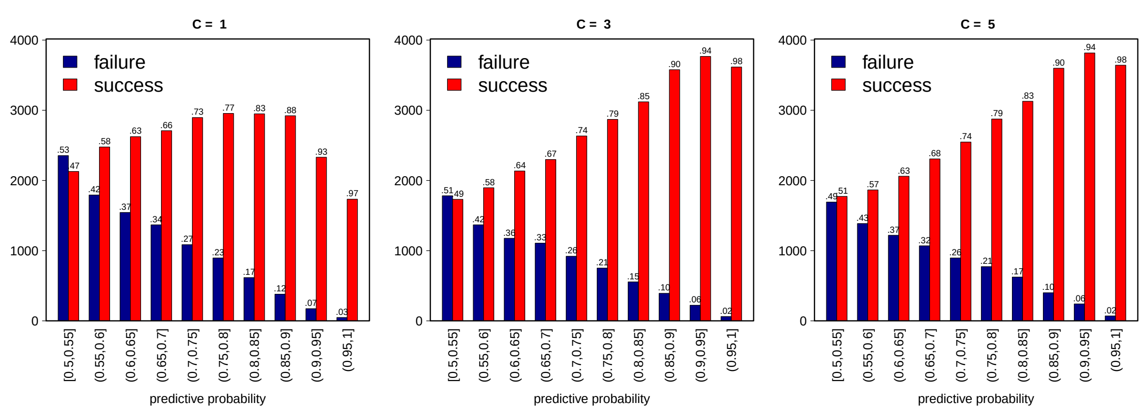

Figure 5: Results of the simulation in Section 4.4. Barplots of the total numbers of successes (red columns) and failures (blue columns) obtained fixing $C = 1$ (left), 3 (middle), and 5 (right), for the data generated with $\lambda_T = 10$. For $C = 1$, 71% of all predictions was correct, for $C = 3$, 76.8%, and for $C = 5$, 76.7%.

References

A. Gelman, G. O. Roberts, and W. R. Gilks. Efficient Metropolis jumping rules. *Bayesian statistics*, 5(599-608):42, 1996.

P. J. Green and X. I. Han. Metropolis methods, gaussian proposals and antithetic variables. In *Stochastic Models, Statistical methods, and Algorithms in Image Analysis*, pages 142–164. Springer, 1992.

E. Irurozki, B. Calvo, and A. Lozano. PerMallows: An R package for Mallows and generalized Mallows models. *Journal of Statistical Software*, 71, 2016.

S. Mukherjee. Estimation in exponential families on permutations. *The Annals of Statistics*, 44(2):853–875, 2016.

G. O. Roberts, A. Gelman, and W. R. Gilks. Weak convergence and optimal scaling of random walk Metropolis algorithms. *The annals of applied probability*, 7(1):110–120, 1997.