A Sequential Monte Carlo Framework for the Mallows Model

Anja Stein†, David Leslie† and Arnoldo Frigessi‡

†STOR-i CDT, Lancaster University ‡Department of Biostatistics, Oslo University

The Mallows Model

Assume we have a set of m distinct items, labelled $\mathcal{A} = \{A_1, \dots, A_m\}$ and we are asked to rank these items in order of our personal preference with respect to some attribute. This information can be expressed as a ranking $\mathbf{R} = (R_1, \dots, R_m)$, where each $R_i \in \{1, \dots, m\}$ corresponds to the rank of an item A_i.

$\mathcal{A} =$ 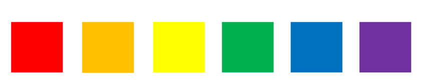

$$\mathbf{R} = (2,\ 4,\ 1,\ 5,\ 3,\ 6)$$

Figure: An example of a ranking for 6 items based on colour preference.

If we consider a probability distribution for rankings, then we can use the Mallows model [2]. The probability of observing a ranking $\mathbf{R}$ is

$$p(\mathbf{R}|\boldsymbol{\rho}, \alpha) = \frac{1}{Z_m(\alpha)} \exp\left\{ -\frac{\alpha}{m} d(\mathbf{R}, \boldsymbol{\rho}) \right\},$$

where: $\boldsymbol{\rho}$, the consensus ranking, is the ranking with the highest probability of being observed; $\alpha > 0$, the scale parameter, controls the variance of the rankings around the consensus ranking; the right-invariant distance function, $d(\cdot, \cdot)$, measures the "closeness" of a permutation to the consensus ranking; and $Z_m(\alpha)$ is the normalisation constant, which is independent of the consensus ranking.

Why use Sequential Monte Carlo (SMC)?

Currently, there is a hierarchical Bayesian model for the Mallows model [3] and an associated Metropolis-Hastings-based MCMC algorithm for obtaining a sequence of random samples from the model. This model can be applied to a wide range of preference data.

Our aim is to estimate the parameters of the Mallows distribution $\boldsymbol{\theta} = \{\boldsymbol{\rho}, \alpha\}$ sequentially as we receive new observations from the distribution. Consider a Mallows model, denoted by π, that we are able to receive new observations at times $t = 1, 2, \dots, T$. We want to estimate the sequence of posterior distributions from M rankings at each time step t, i.e.,

$$\pi_t(\boldsymbol{\theta}) = \pi(\boldsymbol{\theta}|\mathbf{R}_{1:M}) \propto p(\boldsymbol{\theta}) p(\mathbf{R}_{1:M}|\boldsymbol{\theta}).$$

However, we would need to run the MCMC algorithm to convergence each time we receive new observations and this can be computationally costly. Instead, we can use SMC algorithms to approximate the sequence of target distributions using a collection of N weighted random samples, termed particles. At each time step, they evolve according to importance sampling and resampling steps as we receive more observations [1]. SMC admits a fixed computational complexity at each time step.

SMC for Mallows Algorithm

Initialise $(t = 0)$

- **Simulate** N particles from the priors $p(\boldsymbol{\theta})$.

Loop for $t = 1, \dots, T$ (p new rankings arrive each time)

- **Reweight** particles from time $t-1$ to t to account for the new observations. Assuming that the observations are independent, we can calculate the incremental weight of each particle as
$$\frac{\pi(\boldsymbol{\theta}|\mathbf{R}_{1:M+p})}{\pi(\boldsymbol{\theta}|\mathbf{R}_{1:M})} = (Z_m(\alpha_{t-1}^{(i)}))^{-M} \exp\left\{ -\frac{\alpha_{t-1}^{(i)}}{m} \sum_{j=M+1}^{M+p} d(\mathbf{R}_j, \boldsymbol{\rho}_{t-1}^{(i)}) \right\},$$
where $\alpha_{t-1}^{(i)}$ and $\boldsymbol{\rho}_{t-1}^{(i)}$ $(i = 1, \dots, N)$ are the current estimated parameter values of the Mallows model, and M is the number of previously observed rankings. We normalise the weights.
- **Resample** particles (with replacement) and set the weights to $\frac{1}{N}$.
- **Move** particles using a MCMC kernel to create particle diversity.

Experiment with Full Rankings

We compared the MCMC algorithm described in [3] and the SMC algorithm by generating $M_{\text{total}} = 100$ rankings on $m = 20$ items from a Mallows model with $\boldsymbol{\rho} = (1, 2, \dots, 20)$ and $\alpha = 2$. We ran the MCMC algorithm for 1000 iterations and discarded the first 500 iterations as burn-in. In the SMC algorithm, we generated 500 particles using the Uniform prior for $\boldsymbol{\rho}$ and a Exponential prior for α. Then, we let the particles evolve over time by introducing $p = 2$ of our generated rankings over 50 times steps.

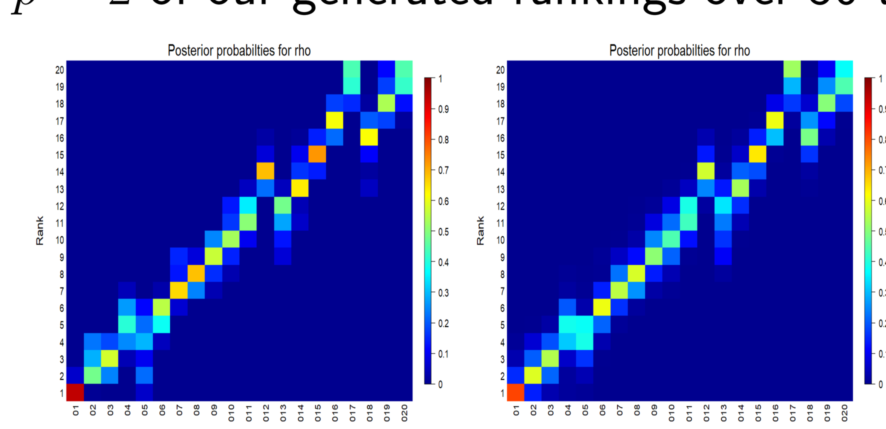

Figure: Heat plot of the posterior probabilities, for 20 items being ranked as the k^{th} most preferred, for $k = 1, \dots, 20$. The plots are the posterior probabilities for $\boldsymbol{\rho}$ using the MCMC algorithm (left) and the SMC algorithm (right). On the x-axis the items are ordered according to their true ranking. The SMC heat plot gives similar posterior probabilities as the MCMC heat plot.

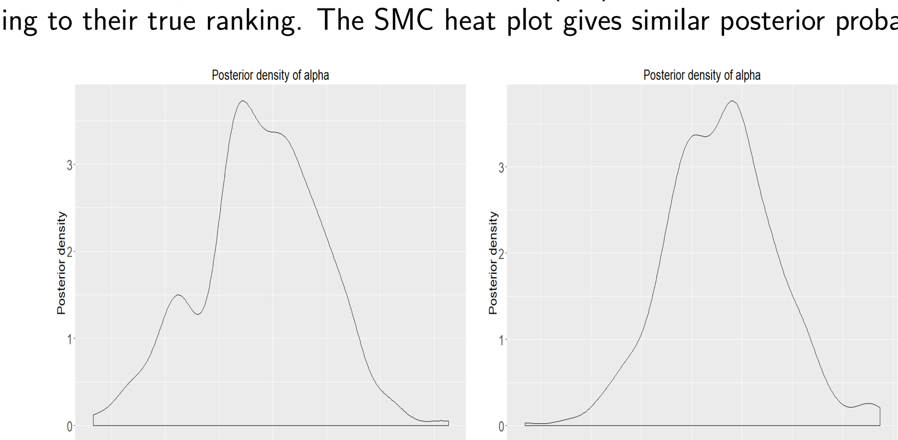

Figure: Posterior density plots for α using the MCMC algorithm (left) and the SMC algorithm (right). Both plots show similar density estimates for α.

Experiment with Partial Rankings

Here, we used the same dataset, but we only observed the top-10 ranked items in each ranking. In the MCMC and SMC algorithms, we needed to use data augmentation to create complete rankings before we were able to estimate the posterior parameters.

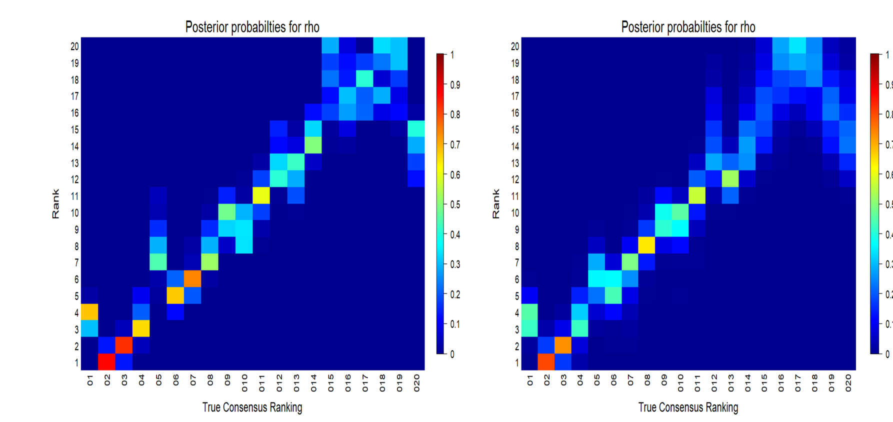

Figure: Heat plot of the posterior probabilities, for 20 items being ranked as the k^{th} most preferred, for $k = 1, \dots, 20$. The plots are the posterior probabilities for $\boldsymbol{\rho}$ using the MCMC algorithm (left) and the SMC algorithm (right). On the x-axis the items are ordered according to their true ranking. The SMC heat plot gives similar posterior probabilities as the MCMC heat plot. It can be seen that there is greater uncertainty for the ranks of the least preferred items.

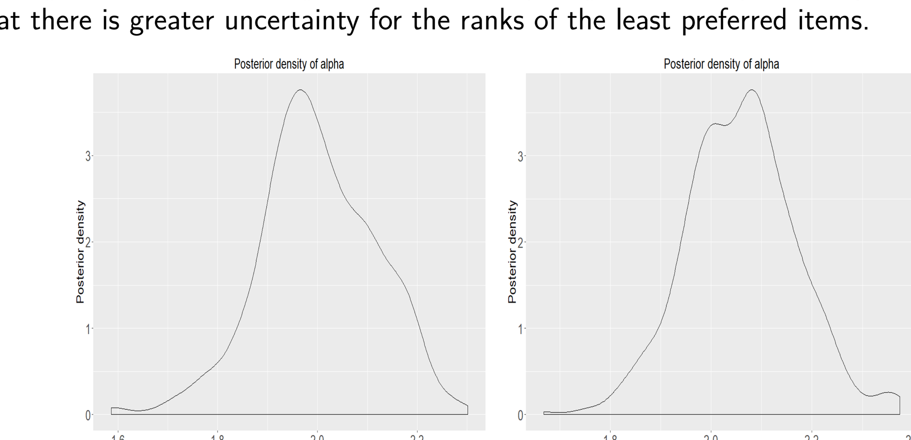

Figure: Posterior density plots for α using the MCMC algorithm (left) and the SMC algorithm (right). Both plots show similar density estimates for α.

Future Work

- Extend the SMC model to the online streaming data scenario where we receive additional information (e.g. an item rank) about an existing partial ranking. It is likely that our augmented rankings will be inconsistent with the new observed ranking. Hence, we need to create an informative kernel to make these corrections based on our current parameter estimates.
- Extend the SMC framework to clusters, where we wish to group a collection of rankings into several groups, each with their own unique parameter values.

References

1. Berzuini Gilks (2001), Resample-Move filtering with cross-model jumps in Sequential Monte Carlo Methods in Practice, *Springer*, 117–138.
2. Mallows (1957), Non-null Ranking Models, *Biometrika*, 44:114-130.
3. Vitelli at al. (2018), Probabilistic Preference learning with the Mallows Rank Model, *Journal of Machine Learning Research*, 18:1–49.

http://www.lancaster.ac.uk/~steina1/

a.stein1@lancaster.ac.uk