

# Rank-based Bayesian variable selection for genome-wide transcriptomic analyses

Emilie Eliseussen<sup>1</sup>  | Thomas Fleischer<sup>2</sup> | Valeria Vitelli<sup>1</sup>

<sup>1</sup>Oslo Centre for Biostatistics and Epidemiology, Department of Biostatistics, University of Oslo, Oslo, Norway

<sup>2</sup>Department of Cancer Genetics, Institute for Cancer Research, Oslo University Hospital, Oslo, Norway

## Correspondence

Valeria Vitelli, Oslo Centre for Biostatistics and Epidemiology, Department of Biostatistics, University of Oslo, P.O. Box 1122, Blindern, 0317 Oslo, Norway.

Email: [valeria.vitelli@medisin.uio.no](mailto:valeria.vitelli@medisin.uio.no)

Variable selection is crucial in high-dimensional omics-based analyses, since it is biologically reasonable to assume only a subset of non-noisy features contributes to the data structures. However, the task is particularly hard in an unsupervised setting, and a priori ad hoc variable selection is still a very frequent approach, despite the evident drawbacks and lack of reproducibility. We propose a Bayesian variable selection approach for rank-based unsupervised transcriptomic analysis. Making use of data rankings instead of the actual continuous measurements increases the robustness of conclusions when compared to classical statistical methods, and embedding variable selection into the inferential tasks allows complete reproducibility. Specifically, we develop a novel extension of the Bayesian Mallows model for variable selection that allows for a full probabilistic analysis, leading to coherent quantification of uncertainties. Simulation studies demonstrate the versatility and robustness of the proposed method in a variety of scenarios, as well as its superiority with respect to several competitors when varying the data dimension or data generating process. We use the novel approach to analyze genome-wide RNAseq gene expression data from ovarian cancer patients: several genes that affect cancer development are correctly detected in a completely unsupervised fashion, showing the usefulness of the method in the context of signature discovery for cancer genomics. Moreover, the possibility to also perform uncertainty quantification plays a key role in the subsequent biological investigation.

## KEYWORDS

Bayesian inference, high-dimensional data, Mallows ranking model, unsupervised learning, variable selection

## 1 | INTRODUCTION

Over the last decades, technological advances have generated an explosion of data in a variety of fields. High-throughput technologies have made -omics data more accessible to researchers, allowing for a systematic exploration of the genetic and epigenetic basis of cancer. However, -omics data are characterized by high-dimensionality, high complexity, and non-normality, and thus pose several significant analytic challenges to most standard statistical approaches. To overcome these complications in high-dimensional omics-based analyses, it is considered biologically reasonable to assume

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2022 The Authors. *Statistics in Medicine* published by John Wiley & Sons Ltd.

that only a small piece of information is relevant for prediction or subtyping,<sup>1-3</sup> so that a low-dimensional solution can also lead to better results interpretability. However, a frequent approach to variable selection for -omics applications consists in performing a priori ad hoc selection of genes according to in-sample statistics or outer literature information, a strategy that heavily affects the solidity and reproducibility of results.<sup>4</sup> Moreover, it has been established in the last decades, with the explosion of high-dimensional data methods, that variable selection should be performed jointly with the main inferential tasks, to properly estimate the uncertainties and the impact of selection on inference results.<sup>5</sup> Among unsupervised methods for variable selection in the context of -omics data, Bayesian model-based approaches are highly suitable, due to their capability of propagating all uncertainties. While heuristic methods may lead to inaccurate estimates because of the low signal-to-noise ratio,<sup>6</sup> especially for high-throughput -omics data, the Bayesian approach allows for incorporating prior biological knowledge when estimating relevant genes from the data. Current methods for Bayesian variable selection for high-dimensional unsupervised problems include probabilistic sparse principal component analysis,<sup>7,8</sup> and Bayesian graphical structure learning,<sup>9,10</sup> all with applications to -omics data. In the case of ranking data, there currently exists no Bayesian unsupervised methods capable of handling the typical dimensions of -omics data.

Converting continuous data to rankings and then performing rank-based analyses is a practice that has gained much popularity in statistical genomics.<sup>11</sup> Different -omics data layers can have incomparable absolute scales, and the use of ranks leads to robust inference.<sup>12</sup> Moreover, analyzing rankings instead of continuous variables allows for easy integration of multiple heterogeneous data, and enhances the reproducibility and interpretability of results. The analysis of rank and preference data has recently seen a growing interest in the machine learning community<sup>13</sup> with the Plackett-Luce<sup>14,15</sup> and Mallows model<sup>16</sup> being among the most commonly used models. The Mallows model is an exponential family-type model on the permutation space, where the probability density depends on a distance between rankings. Compared to Plackett-Luce, the Mallows model shows greater flexibility in the choice of the distance between permutations, and it is also more versatile in adapting to different kinds of data. As a probabilistic model for rank data, the Mallows model enjoys great interpretability, model compactness, inference and computational efficiency. Our current work is based on the Bayesian Mallows model (BMM)<sup>17</sup> with implementing solution `BayesMallows`.<sup>18</sup> `BayesMallows` already provides a computationally feasible inferential approach for the most important choices of distance among permutations, and it shows good accuracy when compared to competitors on datasets of moderate size.<sup>19</sup> However, for this method to be applicable to the typical data dimensions in -omics applications, variable selection is crucial. Moreover, from a purely modeling perspective, imposing the Mallows data generating process on the full (genome-wide) ranked list of genes seems highly unrealistic: RNAseq measurements of moderately expressed genes tend to be very noisy, thus making a low-dimensional solution more suited to this kind of data. Imposing a ranking model only on a selection of genes allows focusing on consistently highly expressed genes, since those will be selected as highly ranked items, thus allowing to detect genome-wide high-quality signals in a more robust way. More generally, this leads to an unsupervised variable selection procedure that can learn from the available rankings to distinguish relevant items from background ones in a high-dimensional setting, as well as provide an estimate of the position in the ranking of the selected relevant items.

The aim of this article is to develop a novel extension of the BMM, the lower-dimensional BMM (lowBMM), providing a low-dimensional solution for rank-based unsupervised variable selection in high-dimensional settings. To the best of our knowledge, the only proposal similar in scope to ours is the partition-Mallows model (PAMA)<sup>20</sup> for rank aggregation, which can distinguish relevant items as well as ranking the relevant ones. This method can handle partial lists, incorporate any available covariate information and can estimate the quality of rankers, as it is based on BARD,<sup>11</sup> which assigns aggregated ranks based on the posterior probability that a certain item is relevant. However, PAMA only allows for the Kendall distance in the Mallows model, so that the partition function has a closed form, while the BMM can use any distance. Additionally, PAMA is unable to perform individual-level inference, a typical advantage of the BMM, and even if it is tested on a relatively complete selection of simulated scenarios, large-data examples are lacking. Concerning rank-based variable selection in -omics analyses, the recent work by Cui et al<sup>21</sup> converts microbiome data to rankings and develops a Bayesian variable selection model using ranks to identify the best covariates for modeling the outcome. However, this latter method is fully supervised, and moreover, in the same way as PAMA, it lacks applicability to high-dimensional settings: all data examples are of much smaller dimensions compared to the typical dimension of -omics data ( $> 10^4$ ). Other Bayesian implementations to analyze rank data have been proposed to reduce computational burden,<sup>22-24</sup> to merge rankings from multiple sources,<sup>11,25</sup> and to combine rank data with other types of data.<sup>26</sup> There exist other methods that are not directly related to lowBMM, but can handle rank data, such as the Markov chain-based methods proposed by Lin and colleagues,<sup>27,28</sup> and various non-Bayesian Mallows based methods.<sup>29,30</sup> Our simulations show that the

proposed lowBMM approach is superior to existing methods in high dimensions, and has no competitor in *ultra*-high dimensions.

The remainder of this article is laid out as follows: In Section 2, we describe the novel lowBMM for variable selection, and its computational implementation. Several simulation studies are described in Section 3, aimed at demonstrating the versatility and robustness of the novel method under different scenarios. Included here is also a sensitivity study on the tuning parameters involved in the model, and a comparison with existing methods. In Section 4, we present an application of this modeling approach to RNAseq gene expression data from ovarian cancer patients. We conclude with a brief summary and discussion in Section 5.

## 2 | LOWER-DIMENSIONAL BAYESIAN MALLOWS MODEL

We briefly present the BMM for complete data as introduced in Vitelli et al<sup>17</sup> in Section 2.1, and outline its extension for unsupervised variable selection, lowBMM, in Section 2.2. A description of the inferential procedure can be found in Section 2.3, while Section 2.4 describes the off-line estimation of the scale parameter  $\alpha$ .

### 2.1 | BMM for complete data

Consider a finite set of  $n$  items denoted as  $\mathcal{A} = \{A_1, A_2, \dots, A_n\}$ . A complete ranking is a mapping  $\mathbf{R} : \mathcal{A} \rightarrow \mathcal{P}_n$  that assigns a rank  $R_i \in \{1, \dots, n\}$  to each item in  $\mathcal{A}$  according to some specified feature,\* with the item ranked 1 being the most preferred (ie, showing the largest value of the feature), and the item ranked  $n$  being the least preferred (ie, showing the smallest value of the feature). Thus, a complete ranking  $\mathbf{R} = (R_1, \dots, R_n)$  lies within the space of  $n$ -dimensional permutations  $\mathcal{P}_n$ . Let us assume that  $N$  assessors have provided complete rankings of the  $n$  items,  $\mathbf{R}_j = \{R_{1j}, R_{2j}, \dots, R_{nj}\}$ ,  $j = 1, \dots, N$ . The Mallows model<sup>16</sup> is a probabilistic model for complete rankings  $\mathbf{R} \in \mathcal{P}_n$ , taking the form

$$P(\mathbf{R}|\alpha, \rho) = \frac{1}{Z_n(\alpha, \rho)} \exp \left\{ -\frac{\alpha}{n} d_n(\mathbf{R}, \rho) \right\} 1_{\mathcal{P}_n}(\mathbf{R}), \quad (1)$$

where  $\alpha > 0$  is a positive scale parameter,  $\rho \in \mathcal{P}_n$  is the latent consensus ranking,  $Z_n(\alpha, \rho)$  is the model partition function,  $1_S(\cdot)$  is the indicator function of the set  $S$ , and finally  $d_n(\cdot, \cdot) : \mathcal{P}_n \times \mathcal{P}_n \rightarrow [0, +\infty)$  is a right-invariant<sup>†</sup> distance function between two rankings. We have here chosen to let the distance function explicitly depend on  $n$ , so that  $d_n(\cdot, \cdot)$  denotes the distance computed on the complete set of items.

Several possibilities for choosing this distance function exist, such as the footrule distance, the Spearman distance, and the Kendall distance.<sup>31</sup> In this article, we choose to use the footrule distance, defined as  $d_n(\mathbf{R}, \rho) = \sum_{i=1}^n |R_i - \rho_i|$ , the equivalent of an  $\ell^1$  measure between rankings. The choice of this distance is motivated by its greater computational efficiency compared to its competitors, such as the Kendall distance, which is more computationally intensive.<sup>18</sup> Moreover, the popularity of the Kendall lies in the fact that a closed form of the model partition function  $Z_n(\alpha, \rho)$  exists in this case; however, we have previously proposed computational solutions to be able to use other distances,<sup>18</sup> while also showing that the Kendall is often less accurate.<sup>19</sup> The Spearman distance is also problematic in our case, as it tends to largely penalize extreme data, being the equivalent of an  $\ell^2$  measure between rankings. On the other hand, the equivalent of the  $\ell^1$  between continuous data seems more suited to -omics applications.

Since the footrule distance is right-invariant, the partition function  $Z_n(\alpha, \rho)$  is independent of  $\rho$  and only dependent on  $\alpha$ . Moreover, we assume in the remainder of this article that  $\alpha$  is fixed and known, and we provide strategies to tune this parameter off-line. This simplification is done to avoid having to compute the normalizing constant  $Z_n(\alpha)$  at each iteration of the Markov chain Monte Carlo (MCMC) algorithm (see Section 2.3), thus greatly lightening the computational burden of the algorithm, at the cost of the tuning of a single scalar parameter, which proves to be straightforwardly carried out in practice (see Section 2.4 for further details). Nonetheless, the performance of the approximated approaches to estimate the normalizing constant  $Z_n(\alpha)$  introduced in Vitelli et al<sup>17</sup> greatly depends on the size of  $n$ , with the performance degrading for decreasing  $n$ , and with the approximation becoming increasingly computationally demanding when  $n \rightarrow \infty$ . Thus, fixing  $\alpha$  simplifies the inferential procedure quite substantially.

Given the model in (1) and the choices specified above, the likelihood associated to the observed rankings  $\mathbf{R}_1, \dots, \mathbf{R}_N$  under the Mallows model takes the form:

$$P(\mathbf{R}_1, \dots, \mathbf{R}_N | \rho) = \frac{1}{Z_n(\alpha)^N} \exp \left\{ -\frac{\alpha}{n} \sum_{j=1}^N d_n(\mathbf{R}_j, \rho) \right\} \prod_{j=1}^N 1_{\mathcal{P}_n}(\mathbf{R}_j).$$

In order to use the Bayesian version of this model,<sup>17</sup> we need to specify a prior for  $\rho$ . To stay in an unsupervised setting, we set a uniform prior  $\pi(\rho) = \frac{1}{n!} 1_{\mathcal{P}_n}(\rho)$  in the space  $\mathcal{P}_n$  of  $n$ -dimensional permutations. The posterior distribution for  $\rho$  then becomes

$$P(\rho | \mathbf{R}_1, \dots, \mathbf{R}_N) \propto \pi(\rho) \exp \left\{ -\frac{\alpha}{n} \sum_{j=1}^N d_n(\mathbf{R}_j, \rho) \right\}.$$

To perform inference, Vitelli et al<sup>17</sup> proposed a MCMC algorithm based on a Metropolis-Hastings (MH) scheme (see Sørensen et al<sup>18</sup> for details on the implementation).

## 2.2 | lowBMM for variable selection

We here introduce a generalization of the BMM presented in Section 2.1 to the scopes of variable selection, which in the context of models for rankings assumes the form of a selection of the *relevant* items, that is, the items worth being ranked. We are thus thinking of generalizing BMM to make it suitable to situations in which  $n$  is large or *ultra-large*, and therefore it becomes unrealistic to assume that all  $n$  items can be completely ranked. We define  $\mathcal{A}^* = \{A_{i_1}, \dots, A_{i_{n^*}}\}$  as an  $n^*$ -dimensional subset of the original set of items  $\mathcal{A}$ , with  $n^* \ll n$  and  $\mathcal{A}^* \subset \mathcal{A}$ . When defining the lowBMM, the underlying assumption is that only a portion of the data follows the Mallows distribution, while the rest is unranked: this is a very realistic assumption, especially in the large  $n$  setting. We then formally assume that items in  $\mathcal{A}^*$  are such that  $R_j|_{\mathcal{A}^*}, j = 1, \dots, N$ , follows a Mallows model over the permutation space of dimension  $n^*$ , where with the notation  $R_j|_S$  we indicate the restriction of an  $n$ -dimensional ranking to only the items belonging to the set  $S$ , which is also a ranking in dimension  $|S|$ . The remaining items in  $\mathcal{A} \setminus \mathcal{A}^*$  are simply irrelevant to the scopes of the analysis, and therefore show no specific pattern in the data. This is equivalent to assuming that  $R_j|_{\mathcal{A} \setminus \mathcal{A}^*}, j = 1, \dots, N$ , is uniformly distributed over the space of permutations of dimension  $n - n^*$ .

$\mathcal{A}^*$  is then the novel model parameter characterizing the lowBMM, the variable selection setting for the BMM. Not only, lowBMM is assuming the data follow a Mallows model only on the lower-dimensional set of items, so that the consensus ranking parameter is now defined as  $\rho \in \mathcal{P}_{n^*}$ . Therefore, the likelihood of lowBMM is defined on a lower-dimensional space of dimension  $n^*$ , and it takes the form:

$$P(\mathbf{R}_1, \dots, \mathbf{R}_N | \rho, \mathcal{A}^*) = \frac{1}{Z_{n^*}(\alpha)^N} \exp \left\{ -\frac{\alpha}{n^*} \sum_{j=1}^N d_{\mathcal{A}^*}(\mathbf{R}_j, \rho) \right\} \prod_{j=1}^N \{U_{\mathcal{P}_{n-n^*}}(\mathbf{R}_j|_{\mathcal{A} \setminus \mathcal{A}^*}) 1_{\mathcal{P}_{n^*}}(\mathbf{R}_j)\},$$

where  $d_{\mathcal{A}^*}(\mathbf{R}_j, \rho) := d_{n^*}(\mathbf{R}_j|_{\mathcal{A}^*}, \rho)$ . This is the same distance as  $d_n(\cdot, \cdot)$  but restricted to the  $n^*$ -dimensional set of items included in  $\mathcal{A}^*$ . The vector  $\mathbf{R}_j|_{\mathcal{A}^*}$  is the restriction of  $\mathbf{R}_j$  to the set  $\mathcal{A}^*$ : it thus defines a partial ordering from which an  $n^*$ -dimensional permutation can be derived. Moreover,  $U_S$  is the uniform distribution over the domain  $S$ , meaning that the term  $U_{\mathcal{P}_{n-n^*}}(\mathbf{R}_j|_{\mathcal{A} \setminus \mathcal{A}^*})$  refers to the assumption that  $\mathbf{R}_j|_{\mathcal{A} \setminus \mathcal{A}^*}$  includes only noisy unranked data.

*Remark 1.* The way the modeling framework is defined does not result in a variable selection procedure that selects the top-ranked items only. Indeed, the items in  $\mathcal{A}^*$  do not need to be those that are most often top- $n^*$  ranked in the data, even if this is the most intuitive solution to a lower-dimensional ranking problem. The relevant items could also show a “consistency” pattern, that is, be often ranked in the same respective order across assessors: this latter possibility constitutes a valid data generating process for lowBMM, as proved in the simulation study in Section 3.1, as the only requirement for lowBMM to be a suitable model for the data is that the ranks of the items follow a Mallows model in the lower-dimensional space  $\mathcal{P}_{n^*}$ . In other words, the model is able to pick out the  $n^*$  items on which the

assessors agree the most, as the model assigns the largest probability to the items whose ranks show the smallest distance to the consensus, which is constrained to be an  $n^*$ -dimensional permutation. In conclusion, the results from lowBMM will depend on the specific dataset: if there exists a top-rank solution in the data, such a solution will be estimated by lowBMM as the one with the largest posterior. On the other hand, if there exists a pattern in the data showing certain items consistently ranked in a specific order (not necessarily top-ranked), lowBMM will give those items the largest probability in the marginal posterior distribution of  $\mathcal{A}^*$ , and rank them accordingly in the marginal posterior distribution of  $\rho$ .

Since the inferential approach is Bayesian, we have to decide priors for all parameters. As we did previously, we set a uniform prior on  $\rho$ , however restricted to the  $n^*$ -dimensional space of permutations of elements of  $\mathcal{A}^*$ :  $\pi(\rho|\mathcal{A}^*) = \frac{1}{n^*!} \mathbf{1}_{\mathcal{P}_{n^*}}(\rho)$ . We also set a uniform prior for  $\mathcal{A}^*$  over  $\mathcal{C}$ ,  $\pi(\mathcal{A}^*) = \frac{1}{|\mathcal{C}|} \mathbf{1}_{\mathcal{C}}(\mathcal{A}^*)$ , where we define  $\mathcal{C}$  as the collection of all  $\binom{n}{n^*}$  possible sets of items of dimension  $n^*$  chosen from a set of dimension  $n$ . The posterior distribution of the variable selection for the Bayes Mallows model can then be written as

$$P(\rho, \mathcal{A}^* | \mathbf{R}_1, \dots, \mathbf{R}_N) \propto \pi(\mathcal{A}^*) \pi(\rho | \mathcal{A}^*) \frac{1}{Z_{n^*}(\alpha)^N} \exp \left\{ -\frac{\alpha}{n^*} \sum_{j=1}^N d_{\mathcal{A}^*}(\mathbf{R}_j, \rho) \right\} \\ \times \prod_{j=1}^N \{ U_{\mathcal{P}_{n-n^*}}(R_j |_{\mathcal{A} \setminus \mathcal{A}^*}) \mathbf{1}_{\mathcal{P}_n}(\mathbf{R}_j) \}, \quad (2)$$

which, by removing the terms not depending on any of the model parameters and by assuming the data are proper rankings, can be simplified to

$$P(\rho, \mathcal{A}^* | \mathbf{R}_1, \dots, \mathbf{R}_N) \propto \exp \left\{ -\frac{\alpha}{n^*} \sum_{j=1}^N d_{\mathcal{A}^*}(\mathbf{R}_j, \rho) \right\} \mathbf{1}_{\mathcal{P}_{n^*}}(\rho) \mathbf{1}_{\mathcal{C}}(\mathcal{A}^*). \quad (3)$$

Note that, from the posterior distribution in (3), marginal posterior distributions of both model parameters,  $\rho$  and  $\mathcal{A}^*$ , can be easily derived. Often one is interested in computing posterior summaries of such distributions. By inspecting the marginal posterior distribution of  $\mathcal{A}^*$  estimated by lowBMM, one can for instance a posteriori select the items in  $\mathcal{A}$  that maximize this marginal posterior under all possible model reductions of dimension  $n^*$ , thus practically attaining variable selection via the marginal posterior mode (ie, the maximum a posteriori—MAP) of  $\mathcal{A}^*$ . However, computing the MAP might not be the best way of summarizing the marginal posterior distributions of both  $\rho$  and  $\mathcal{A}^*$ , in addition to being a computationally intensive procedure for larger  $n$  and  $n^*$ . Instead, we introduce novel ways for providing posterior summaries of both the consensus ranking of the relevant items, called  $\hat{\rho}_{\mathcal{A}^*}$ , and of the set of relevant items, called  $\hat{\mathcal{A}}^*$ ; these approaches are described in Section 2.3.1.

Furthermore, the inspection of the marginal posterior distributions of  $\rho$  and  $\mathcal{A}^*$  has much wider implications than simply obtaining posterior summaries: for example, by inspecting the marginal posterior distribution of  $\mathcal{A}^*$  one can decide whether variable selection is appropriate to the data at hand (few items are clearly assigned rankings while the rest shows high uncertainty), or rather items are a posteriori ranked in blocks (several items share the same rank with large probability, and only a respective ordering between groups of items can be properly estimated). Scope of the simulation studies in Section 3 is to assess the correct behavior of the model in these different scenarios, together with its gain in efficiency in the large  $n$  context.

*Remark 2.* Note that we have assumed  $n^*$  to be fixed and known. Estimating  $n^*$  together with the other model parameters would imply a model selection step, which would make the inferential task substantially more challenging, or require to completely change our inferential approach, for instance toward the use of stick-breaking priors.<sup>32</sup> Both possibilities are completely out of the scope of the present article and are left to future speculations. On the other hand, the choice of fixing  $n^*$  is typically not problematic in practice, as one can often tune it in connection to the real application (eg, in genomics, biologists are often interested in a specific “number” of relevant genes, even more if we consider that each gene-regulating biological pathway includes on average a hundred genes<sup>33</sup>). Another practical possibility is to simply choose  $n^*$  as large as allowed by the computing resources at hand. This latter approach is also supported by the simulation studies (see Section 3.2), where we realized that choosing  $n^*$  larger than necessary would result in a solution including and automatically detecting the lower dimensional relevant subset.



## 2.3 | MH algorithm for inference in the lowBMM

In order to obtain samples from the posterior in Equation (2), we set up a MH-MCMC scheme. The algorithm iterates between two steps: in one step the consensus ranking  $\rho$  given the current set  $\mathcal{A}^*$  is updated, and in the other the set  $\mathcal{A}^*$  is updated given the current consensus ranking.

In the first step of the algorithm, we propose a new consensus ranking  $\rho' \in \mathcal{P}_{n^*}$  using the “leap-and-shift” proposal distribution.<sup>17</sup> The acceptance probability for updating  $\rho$  in the MH algorithm is

$$\min \left\{ 1, \frac{P_l(\rho|\rho')}{P_l(\rho'|\rho)} \exp \left[ -\frac{\alpha}{n^*} \left( \sum_{j=1}^N d_{\mathcal{A}^*}(\mathbf{R}_j, \rho') - \sum_{j=1}^N d_{\mathcal{A}^*}(\mathbf{R}_j, \rho) \right) \right] \right\}. \quad (4)$$

In Equation (4),  $P_l$  denotes the probability mass function associated to moving from one  $n^*$ -dimensional ranking vector to another according to the leap-and-shift proposal. The parameter  $l$  denotes the number of items perturbed in the consensus ranking  $\rho$  to get a new proposal  $\rho'$ , and is used to tune the acceptance probability. To compute the distances restricted to  $\mathcal{A}^*$ , the  $\mathbf{R}_j$ 's needs to be updated so that they correspond to the items in the current set  $\mathcal{A}^*$ .

The second step of the algorithm updates the set  $\mathcal{A}^*$ . We propose a new set  $\mathcal{A}_{\text{prop}}^*$  by perturbing  $L \in \{1, \dots, n^*\}$  elements in the current  $\mathcal{A}^*$ , selected with uniform probability. The  $L$  items are swapped with  $L$  items from the set  $\mathcal{A} \setminus \mathcal{A}^*$ , again uniformly. Therefore, we can formally write the proposal distribution as  $q(\mathcal{A}_{\text{prop}}^*|\mathcal{A}^*) = |C_L^{\text{in}}|^{-1} |C_L^{\text{out}}|^{-1} \mathbf{1}_{C_L^{\text{in}}}(\mathcal{A}_{\text{prop}}^* \setminus \mathcal{A}^*) \mathbf{1}_{C_L^{\text{out}}}(\mathcal{A}^* \setminus \mathcal{A}_{\text{prop}}^*)$ , where we define  $C_L^{\text{in}}$  as the collection of all  $\binom{n-n^*}{L}$  possible sets of items of dimension  $L$  chosen from a set of dimension  $n - n^*$ , to be brought “in” to the reduced set  $\mathcal{A}^*$  in order to obtain  $\mathcal{A}_{\text{prop}}^*$ . Likewise, we define  $C_L^{\text{out}}$  as the collection of all  $\binom{n^*}{L}$  possible sets of items of dimension  $L$  chosen from a set of dimension  $n^*$ , to be brought “out” of the reduced set  $\mathcal{A}^*$  in order to obtain  $\mathcal{A}_{\text{prop}}^*$ . Note that such proposal distribution is by definition symmetrical. The move from  $\mathcal{A}^*$  to  $\mathcal{A}_{\text{prop}}^*$  is then accepted with probability:

$$\min \left\{ 1, \frac{\pi(\mathcal{A}_{\text{prop}}^*) P(\mathbf{R}_1, \dots, \mathbf{R}_N | \mathcal{A}_{\text{prop}}^*, \rho) q(\mathcal{A}^* | \mathcal{A}_{\text{prop}}^*)}{\pi(\mathcal{A}^*) P(\mathbf{R}_1, \dots, \mathbf{R}_N | \mathcal{A}^*, \rho) q(\mathcal{A}_{\text{prop}}^* | \mathcal{A}^*)} \right\}.$$

Since the prior for  $\mathcal{A}^*$  is uniform, the prior terms in the ratio simplify, as well as the symmetrical proposal distributions. We are then left with a ratio of likelihoods that can be written as

$$\min \left\{ 1, \exp \left[ -\frac{\alpha}{n^*} \left( \sum_{j=1}^N d_{\mathcal{A}_{\text{prop}}^*}(\mathbf{R}_j, \rho) - \sum_{j=1}^N d_{\mathcal{A}^*}(\mathbf{R}_j, \rho) \right) \right] \right\}. \quad (5)$$

Note that  $d_{\mathcal{A}_{\text{prop}}^*}(\mathbf{R}_j, \rho) = d_{n^*}(\mathbf{R}_j |_{\mathcal{A}_{\text{prop}}^*}, \rho)$ , that is, the distance is computed on the restricted set of the current  $\mathcal{A}_{\text{prop}}^*$ , and  $\mathbf{R}_j$  is also restricted to the same set. Therefore, one needs to decide how to evaluate  $\rho$  on  $\mathcal{A}_{\text{prop}}^*$ . There are multiple ways of doing this, the most simple being to randomly perturb the current  $\rho$  from the first step of the algorithm to accommodate the new  $L$  items proposed for  $\mathcal{A}_{\text{prop}}^*$ . That is, given the current  $\rho$ , we discard the ranks corresponding to the  $L$  items that were left out of  $\mathcal{A}_{\text{prop}}^*$ , and we randomly reassign those same ranks to the new items.

The MH-MCMC algorithm described above is summarized in Algorithm 1.

### 2.3.1 | Postprocessing of the MCMC results

In order to derive posterior summaries of the consensus ranking of the relevant items and of the set of relevant items, named  $\hat{\rho}_{\mathcal{A}^*}$  and  $\hat{\mathcal{A}}^*$ , respectively, we inspect the marginal posterior distribution of  $\rho$  and  $\mathcal{A}^*$  (after burn-in). Suppose  $M$  posterior samples are obtained:  $\{\rho_m, \mathcal{A}_m^*\}_{m=1}^M$  with  $\rho_m = \{\rho_{m1}^m, \dots, \rho_{mn^*}^m\}$  and  $\mathcal{A}_m^* = \{A_{m1}^m, \dots, A_{mn^*}^m\}$ . Given the samples  $\{\mathcal{A}_1^*, \dots, \mathcal{A}_M^*\}$ , let  $W \in \mathbb{R}^{M \times n}$  be such that  $W_{mi} = \mathbf{1}_{\mathcal{A}_m^*}(A_i)$  for each item  $A_i$ ,  $i = 1, \dots, n$ . For the computation of the posterior summaries, we first need the following definitions.

**Definition 1.** Given a vector of real numbers  $(x_1, \dots, x_n) \in \mathbb{R}^n$ , its corresponding rank vector is obtained as follows:

$$\text{rank}(x_1, \dots, x_n) = (r_1, \dots, r_n), \quad \text{such that } r_i = \sum_{j=1}^n \delta(x_i - x_j) \quad \text{for } i = 1, \dots, n,$$

$$\text{where } \delta(x) = \begin{cases} 1, & \text{if } x \geq 0, \\ 0, & \text{if } x < 0. \end{cases}$$

**Algorithm 1.** MH-MCMC scheme for inference in lowBMM

**input:**  $R_1, \dots, R_N, \alpha, d(\cdot, \cdot), l, L, M$   
**output:** posterior distributions of  $\rho$  and  $\mathcal{A}^*$   
**Initialization:** randomly generate  $\rho_0$  and  $\mathcal{A}_0^*$   
**for**  $m \leftarrow 1$  **to**  $M$  **do**  
  **M-H step: update  $\rho$**   
  sample:  $\rho' \sim \text{LS}(\rho_{m-1}, l)$  restricted on  $\mathcal{A}_{m-1}^*$  and  $u \sim \mathcal{U}(0, 1)$   
  compute:  $ratio \leftarrow$  Equation (4) with  $\rho \leftarrow \rho_{m-1}, \mathcal{A}^* \leftarrow \mathcal{A}_{m-1}^*$   
  **if**  $u < ratio$  **then**  
  |  $\rho_m \leftarrow \rho'$   
  **else**  
  |  $\rho_m \leftarrow \rho_{m-1}$   
  **end**  
  **M-H step: update  $\mathcal{A}^*$**   
  sample:  $\mathcal{A}_{\text{prop}}^* \sim q(\mathcal{A}_{\text{prop}}^* | \mathcal{A}_{m-1}^*)$ , and  $u \sim \mathcal{U}(0, 1)$   
  compute:  $ratio \leftarrow$  Equation (5) with  $\rho \leftarrow \rho_m, \mathcal{A}^* \leftarrow \mathcal{A}_{m-1}^*$   
  **if**  $u < ratio$  **then**  
  |  $\mathcal{A}_m^* \leftarrow \mathcal{A}_{\text{prop}}^*$   
  **else**  
  |  $\mathcal{A}_m^* \leftarrow \mathcal{A}_{m-1}^*$   
  **end**  
**end**

**Definition 2.** For  $n^* \leq k \leq n$ , the highest probability set (HPS) of  $\mathcal{A}^*$  is defined as follows:

$$\mathcal{A}' = \{A_i, i = 1, \dots, n \mid \text{rank}(\bar{\mathbf{w}})_i \leq k\},$$

where  $\bar{\mathbf{w}} = (\bar{w}_1, \dots, \bar{w}_n)$  and  $\bar{w}_i = \frac{1}{M} \sum_{m=1}^M W_{mi}$ .

Let  $\mathcal{A}'$  be the HPS of  $\mathcal{A}^*$ . Based on  $\mathcal{A}'$  we compute  $\bar{\mathbf{x}} \in \mathbb{R}^{|\mathcal{A}'|}$ ,  $\bar{x}_i = \frac{\sum_{m=1}^M \rho_{mi} 1_{\mathcal{A}_m^*}(A_i)}{\sum_{m=1}^M 1_{\mathcal{A}_m^*}(A_i)}$  for all  $A_i \in \mathcal{A}'$ . We now quantify the two posterior summaries of  $\rho$  and  $\mathcal{A}^*$  as follows:

$$\hat{\mathcal{A}}^* = \{A_i \in \mathcal{A}' \mid \text{rank}(\bar{\mathbf{x}}) \leq n^*\}, \quad \hat{\rho}_{\mathcal{A}^*} = \text{rank}(\bar{\mathbf{x}})|_{\hat{\mathcal{A}}^*}. \quad (6)$$

Note that in all the figures in the article reporting heatplots of the marginal posterior distribution of  $\rho$ , the items on the x-axis will be ordered according to their respective ranking in  $\hat{\rho}_{\mathcal{A}^*}$ .

In some cases, it might be of interest to obtain a “top probability selection within the selection”  $\hat{A}_{\text{top}}^* \in \hat{\mathcal{A}}^*$  for further downstream analysis (for instance in cases when not only  $n$  but also  $n^*$  is quite large). The strategy for computing  $\hat{A}_{\text{top}}^*$  consists in inspecting the probability distributions  $P(A_i \in \text{top-}K, i = 1, \dots, n \mid A_i \in \hat{\mathcal{A}}^*)$  for varying  $K \leq n^*$ . From these probability distributions, we would like to choose a  $K$  that discriminates the items that have a high probability of being selected compared to the rest of the items. Once the  $K$  is specified, the probability distribution will by construction be bimodal, and it will be possible to find a cut-off point  $c$  that discriminates its two peaks well (this has explicitly been explained in the context of the case study presented in Section 4, see Figure 11). The cut-off  $c$  is then used as a lower probability bound for being included in  $\hat{A}_{\text{top}}^*$ . The result is a “top probability selection” of  $|\hat{A}_{\text{top}}^*|$  items where  $\hat{A}_{\text{top}}^* = \{A_i \in \hat{\mathcal{A}}^* \text{ s.t. } P(A_i \in \text{top-}K, i = 1, \dots, n \mid A_i \in \hat{\mathcal{A}}^*) > c\}$ .

## 2.4 | Off-line estimation of $\alpha$

To simplify inference by avoiding the computation of the normalizing constant, we decided to fix  $\alpha$ . Nonetheless, we would still need an estimation of its most reasonable value when we are given a new ranking dataset. We propose one possible method for estimating such a value, inspired by Liu et al.<sup>34</sup>

As  $\alpha$  describes how closely the rankings are distributed around the common consensus, or in other words, how similar the individual rankings are to each other, we can assume that a higher  $\alpha$  indicates a higher “similarity” between assessors. Therefore, when assuming that the individual rankings of the assessors are drawn from a Mallows with parameter  $\alpha$ , a higher  $\alpha$  results in a lower mean distance between the assessors. To estimate  $\alpha$ , we assume that the mean distance among  $N$  assessors for a real ranking dataset should resemble the mean distance of a simulated ranking dataset of the same dimension, generated by drawing independent samples from the Mallows distribution. Precisely, suppose we want to estimate  $\alpha$  for a real ranking dataset with  $n$  items ranked by  $N$  assessors: we can then generate several ranking datasets of dimension  $(N, n)$  over a grid of  $\alpha_0$  values,  $\mathbf{R}_1^{\alpha_0}, \dots, \mathbf{R}_N^{\alpha_0}$  (the chosen consensus ranking  $\rho_0$  does not matter). Then, we calculate the mean pairwise distance between assessors for each simulated dataset (ie, for every  $\alpha_0$ ), defined as:

$$\bar{d}_{\alpha_0} = \frac{1}{N(N-1)} \sum_{j=1}^N \sum_{k \neq j} d(\mathbf{R}_j^{\alpha_0}, \mathbf{R}_k^{\alpha_0}). \quad (7)$$

Finally, we compute the same mean distance between assessors as in (7) for the dataset we are interested in, and choose the  $\alpha_0$  value where the two distance measures intersect.

The approach to estimate  $\alpha$  described so far works for a complete ranking dataset where we assume that all  $n$  items are ranked according to independent draws from the Mallows distribution. However, when using the lowBMM, we will be assuming that only  $n^*$  items are ranked according to a Mallows model. Hence, when using the method described above we will need to rescale the estimated  $\alpha$  from dimension  $(N, n)$  to dimension  $(N, n^*)$ . Let  $\hat{\alpha}_n$  be the optimal  $\alpha_0$  value estimated from the datasets of dimension  $(N, n)$ , as described above. We can then rescale  $\hat{\alpha}_n$  to obtain the optimal  $\hat{\alpha}_{n^*}$  value in dimension  $n^*$  by matching the terms in the exponent in the Mallows model:

$$\hat{\alpha}_{n^*} = \hat{\alpha}_n \frac{n \max_{d_{n^*}}}{n^* \max_{d_n}},$$

where  $\max_{d_n}$  (respectively  $\max_{d_{n^*}}$ ) is the maximum attainable footrule distance between two rankings in dimension  $n$  (respectively  $n^*$ ). The method is assessed in a simulation study described in Section 3.5.

### 3 | SIMULATION EXPERIMENTS

This section describes how lowBMM performs under different data generating processes, how its performance is affected by the various parameters involved in the MCMC estimation, and how it compares to other methods. The variable selection procedure provides a posterior distribution for the consensus ranking in a reduced dimension, and the type of distribution obtained in turn depends on the data generating process. Therefore, inspecting the solution provided by lowBMM under different scenarios for generating the data allows us to understand the data structure and patterns in real-life situations.

First, we consider a top-rank data generating process, where only a subset of the items is ranked, and assigned the top ranks in reduced dimension, while the rest is noisy: this scenario results in a top-rank solution, that is, lowBMM correctly selects the top-items following a rank model. Second, we generate data where items are consistently awarded ranks in a specific order: this results in a rank-consistency solution, that is, items that *can* be ordered are selected and top-ranked. A study on how lowBMM works under different data generating processes is presented in Section 3.1, and a sensitivity analysis on the tuning parameters involved in lowBMM can be found in Section 3.2. We also test the method’s robustness to the noise level in the data in Section 3.3, and we provide an extensive study on how lowBMM compares to existing methods in Section 3.4. Finally, we present a performance assessment of the off-line estimation of  $\alpha$  in Section 3.5.

#### 3.1 | Varying the data generating process: top-rank and rank consistency experiments

To test the method’s applicability to different data structures, we consider two data generating processes: top-rank and rank-consistency. In the case of top-rank we generate data in the following way:  $n^*$  items were given ranks sampled from the Mallows model,  $R_j|_{A^*} \sim \text{Mallows}(\rho_{A^*}, \alpha)$ , with  $\rho_{A^*} = (1, \dots, n^*)$  while the rest of the items were randomly assigned

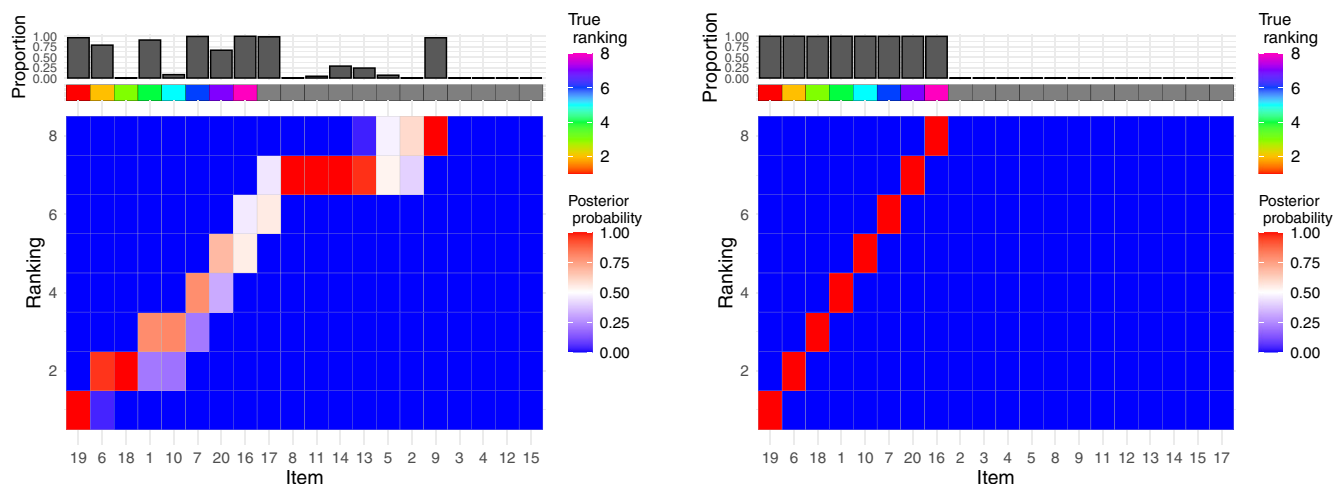


ranks,  $R_j|_{\mathcal{A}\setminus\mathcal{A}^*} \sim \mathcal{U}(\mathcal{P}_{n^*+1, \dots, n})$ . Here the top items are ranked according to the Mallows model, and thus result in a top-rank solution. In the rank consistency experiment on the other hand, the aim is to see whether lowBMM can select consistently ranked items even if not necessarily top-ranked. The setting was the following: items in  $\mathcal{A}^*$  were given a respective ranking sampled from a Mallows model in dimension  $n^*$ ; then, for obtaining each assessor-specific ranking, the same items were assigned a random ranking in dimension  $n$  in their respective positions; finally, the rest of the items had randomly assigned ranks. Therefore, items in the true  $\mathcal{A}^*$  were always ranked in a certain order for each assessor, however their specific ranks were not consistent.

For the top-rank data generating process, we performed two data simulation experiments: first a toy example with  $n = 20$  items, and then a larger example with  $n = 1000$  items.

In the small example, we considered a group of  $N = 50$  assessors, and  $n^* = 8$  relevant items. For what concerns the tuning parameters, we set  $L = 1$  and  $l = \text{round}(n^*/5)$  as suggested by the tuning parameter study in Section 3.2. We used  $M = 5 \times 10^3$  MCMC iterations, and the algorithm took 10 seconds to run. The left panel in Figure 1 displays the marginal posterior distribution of  $\rho$  where the items on the  $x$ -axis have been ordered according to  $\hat{\rho}_{\mathcal{A}^*}$  (see Section 2.3.1 for the computation of posterior summaries). The results show that the method is able to correctly select the relevant items to be included in the selection set, as the items whose ranking was simulated from the Mallows model are all top-ranked. However, there is uncertainty associated to their specific rankings: some items seem to “compete” for ranks, for example, item 6 (true rank 2) and item 18 (true rank 3) both show a large posterior marginal probability of being ranked 2; at the same time, their final ordering is correct (the ordering on the  $x$ -axis is consistent with the rainbow-grid representing the true ordering  $\rho_{\mathcal{A}^*}$ ). This is not an unexpected behavior, as the variability in the items’ observed ranks is dependent on the scale parameter,  $\alpha$ . As we do not estimate  $\alpha$  in the current version of lowBMM, we inspected how the solution changed when varying this parameter. The  $\alpha$  parameter regulates how much the assessors agree with each other on the ranking of the items: a large  $\alpha$  means a better agreement between the assessors, and a lower  $\alpha$  means less agreement. So, when  $\alpha$  is large, the items distributed according to the Mallows are easier to pick out, making it less likely that other items are accepted in the proposal set. This is evident in the distribution of the bars on the top of the plots in Figure S1 in the supplementary material, where the larger  $\alpha$  value shows less variability in the selected items. Moreover, when the data is “easier,” the algorithm converges much more quickly to the correct set. The bottom items (whose ranks are essentially noise) are clearly distributed on the bottom for  $\alpha = 10$ , while there is a bit more uncertainty for  $\alpha = 3$ , as expected.

In order to test the method in a more realistic scenario, we simulated the data in a dimension closer to that often encountered in -omics applications. This larger experiment aims to investigate the method performance and its computing capabilities on a more demanding dataset when using the same data generating process as before (top-rank). We set the total number of items to  $n = 1000$ , the number of relevant items to  $n^* = 50$ , and the number of assessors to  $N = 50$ . We

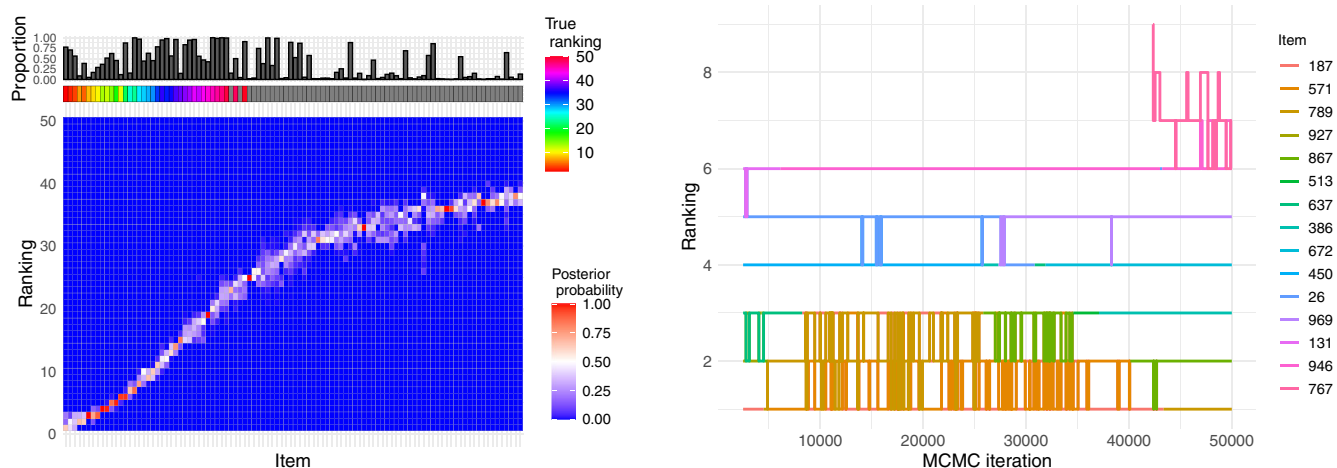


**FIGURE 1** Results from the small simulation experiment described in Section 3.1 with  $\alpha = 10$ ,  $n = 20$ ,  $N = 50$ ,  $n^* = 8$ ,  $L = 1$ , and  $l = \text{round}(n^*/5)$ : Both panels display heatplots of the marginal posterior distribution of  $\rho$  with the items ordered according to  $\hat{\rho}_{\mathcal{A}^*}$  on the  $x$ -axis. Left: Top-rank simulation example, right: rank consistency simulation example. The rainbow grid indicates the true ordering  $\rho_{\mathcal{A}^*}$ , and the bar plot indicates the proportion of times the items were selected in  $\mathcal{A}^*$  over all MCMC iterations

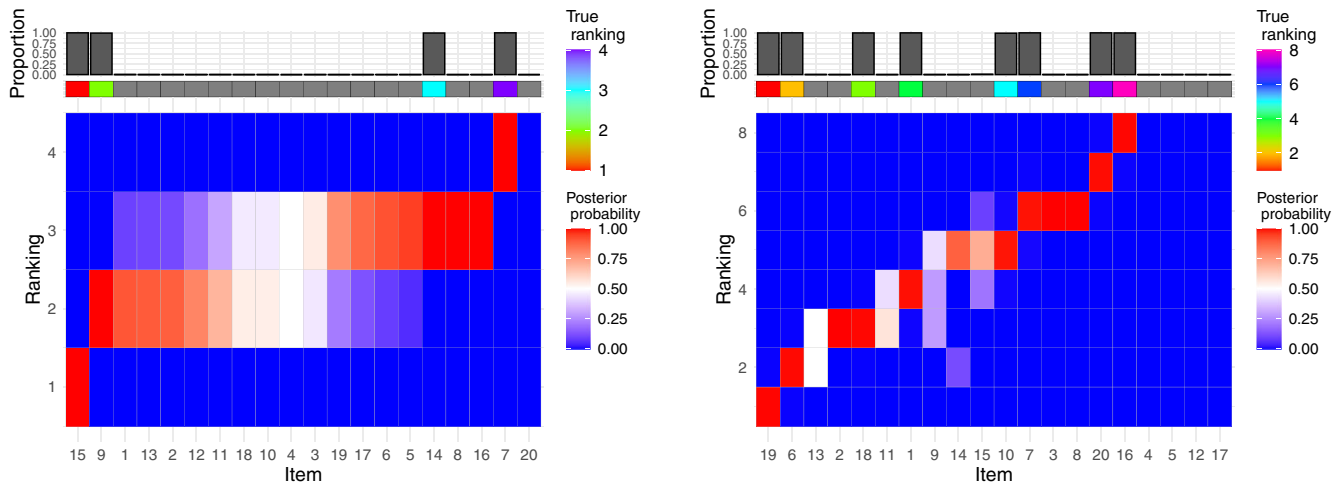
also set the tuning parameters to  $L = 1$  and  $l = \text{round}(n^*/5)$ , as previously done. We ran the MCMC for  $5 \times 10^4$  iterations in this larger experiment, and the algorithm used 155 seconds to run. Note that we used a larger  $M$  in the larger examples to ensure good space exploration and convergence of the chains, and the computing times were all scaling well with the increasing dimension. The posterior of  $\rho$  can be seen in the left plot of Figure 2: more than 75% of the items in the true selection are assigned top ranks with large marginal posterior probability, while the uncertainty associated to the items' rankings increases moving toward the bottom, for the items not in the top- $n^*$ . The trace plot on the right in Figure 2 demonstrates convergence, with most top-ranked items clearly converging to one final rank, and only a few items competing for the same rank throughout the chain. This competing behavior is to be welcomed if the posterior distribution supports it, which seems to be the case here as this demonstrates a higher uncertainty in the rank of the items in question, as is to some degree expected in high-dimensional datasets with a large number of items. We noticed this behavior to be enhanced by increasing the tuning parameter  $L$  (see Figure S9 in the supplementary material), suggesting that in such cases convergence is more difficult to achieve, and therefore we advise to keep  $L$  small (this is confirmed in the sensitivity study on the tuning parameters in Section 3.2).

For the rank consistency data generating process, we carried out a simulation study similar to the small top-rank example, also here focusing on a small experiment with  $n = 20$ ,  $n^* = 8$ , and  $N = 50$ . The right panel in Figure 1 presents the marginal posterior distribution of  $\rho$  where the items have been ordered according to  $\hat{\rho}_{\mathcal{A}^*}$  on the  $x$ -axis. Here the performance of lowBMM is evident: the correct items (indicated by the rainbow grid on the top) are given the highest marginal posterior probability in  $\mathcal{A}^*$  (as shown by the bar plot on top), and the ordering of the items is also consistent with the true  $\rho_{\mathcal{A}^*}$ . This shows that lowBMM is able to handle varying data structures, and quite interestingly, is performing slightly better on a dataset with consistently ranked items, compared to the top-rank data example. The method also seems to be somewhat more robust to the choice of the tuning parameters for this type of data generating process. The rank consistency simulation experiment can be considered a more realistic example of how ranking data could be structured in real-life situations, particularly in genomics: a subset of genes might be consistently ranked in a certain pattern across samples (eg, consistently more or less expressed), while their individual rankings do not really matter as much. At the same time, such a pattern would be much more interesting to detect than the actual true ranking estimation. Furthermore, experiments making use of this data generating process were overall more challenging to handle for all the other competing methods considered for comparison (see Section 3.4).

In the rank consistency simulation experiment, we also experimented by reducing the number of items being given the consistent ranks,  $n^*$ , as well as the level of precision in the data sampled from the Mallows,  $\alpha$ . Decreasing both  $n^*$  and  $\alpha$  makes the estimation via lowBMM more difficult, as less items show a detectable pattern, and the noise level in the data is higher. Figure 3 shows the results obtained with  $\alpha = 0.5$ : as expected, we see a much larger uncertainty in the ranks with respect to what observed in Figure 1 (right) for  $\alpha = 10$ . It is interesting to note the performance of the



**FIGURE 2** Results from the large top-rank simulation experiment described in Section 3.1 with  $\alpha = 10$ ,  $n = 1000$ ,  $N = 50$ ,  $n^* = 50$ ,  $L = 1$ ,  $l = \text{round}(n^*/5)$ . Left: Heatplot of the marginal posterior distribution of  $\rho$  for the top-100 items ordered according to  $\hat{\rho}_{\mathcal{A}^*}$  on the  $x$ -axis, where the rainbow grid on top indicates the true ranking of the relevant items  $\rho_{\mathcal{A}^*}$ , and the bar plot indicates the proportion of times the items were selected in  $\mathcal{A}^*$  over all MCMC iterations. Right: Trace plot of the top-15 items in  $\hat{\rho}_{\mathcal{A}^*}$  along MCMC iterations



**FIGURE 3** Results from the rank consistency simulation experiment described in Section 3.1: Heatplot of the marginal posterior distribution of  $\rho$  with the items ordered according to  $\hat{\rho}_{\mathcal{A}^*}$  on the  $x$ -axis with  $\alpha = 0.5$ ,  $n = 20$ ,  $N = 50$ ,  $L = 1$ , and  $l = \text{round}(n^*/5)$ . Left:  $n^* = 4$ ; right:  $n^* = 8$ . The rainbow grid indicates the true ranking of the relevant items  $\rho_{\mathcal{A}^*}$ , and the bar plot indicates the proportion of times the items were selected in  $\mathcal{A}^*$  over all MCMC iterations


method: despite the much larger uncertainty, still most weight in the marginal posterior of  $\rho$  is given to the correct items (looking at the bar plots at the top), and the ordering of the items is consistent with the true ordering (rainbow grid on top). This shows that the accuracy of lowBMM is unaffected by decreasing  $\alpha$  or  $n^*$ , even in the case of a rank consistency data generating process. Note that in this experiment we have used the same  $n^*$  both in the data generation and in the analysis, as the main point here was assessing the robustness of lowBMM to a less detectable pattern. In Section 3.2, we will instead verify the method robustness to the misspecification of  $n^*$ .

### 3.2 | Sensitivity study on tuning parameters

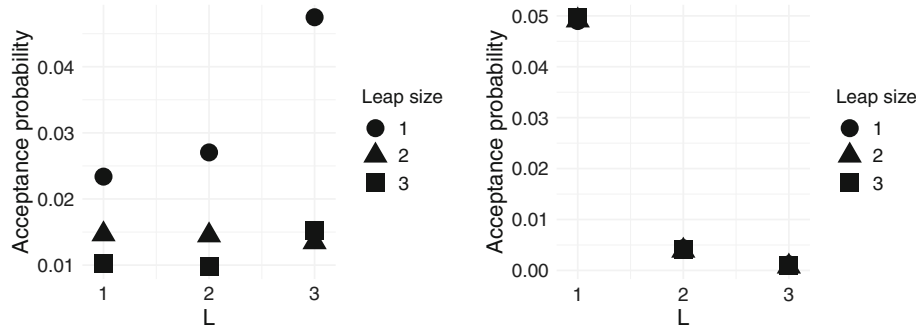
A sensitivity study was carried out to verify the effect of the tuning parameters  $L$  and  $l$  on lowBMM, as well as the effect of misspecifying the number of relevant items for the selection,  $n^*$ . We first focused on studying the effect of  $L$  and  $l$  and therefore we set:  $n = 100$ ,  $N = 10$ ,  $n^* = 10$ , and  $\alpha = 5$ . The data were generated according to the top-rank data generating process described in Section 3.1:  $n^*$  items were given ranks sampled from the Mallows model, while the rest of the items were ranked randomly lower. We experimented by varying the two tuning parameters introduced in Section 2.3:  $L$ , that is, the number of items perturbed in  $\mathcal{A}^*$  for proposing a new set  $\mathcal{A}_{\text{prop}}^*$ , and  $l$ , that is, the number of items perturbed in the current  $\rho$  for proposing a new consensus  $\rho'$ . Values for  $L$  and  $l$  were explored over a grid  $(L, l) \in [1, 2, 3, 4, 5] \times [1, 2, 3, 4, 5]$ . Results were collected after running the method on 20 simulated datasets, when using  $M = 5000$  MCMC iterations, which always showed to be more than enough for convergence.

To evaluate the results, we computed the footrule distance between the true and the estimated consensus ranking, normalized by the number of correctly selected items  $n_{\text{corr}} = |\mathcal{A}^* \cap \hat{\mathcal{A}}^*|$ :  $d_{\text{norm}}(\rho_{\mathcal{A}^*}, \hat{\rho}_{\mathcal{A}^*}) = d(\rho_{\mathcal{A}^*}, \hat{\rho}_{\mathcal{A}^*})/n_{\text{corr}}$  (see Section 2.3.1 for details on the computation of the posterior summaries  $\hat{\rho}_{\mathcal{A}^*}$  and  $\hat{\mathcal{A}}^*$ ). For brevity, we will sometimes indicate  $d_{\text{norm}}(\rho_{\mathcal{A}^*}, \hat{\rho}_{\mathcal{A}^*})$  simply with  $d_{\text{norm}}$ .

The left panel in Figure 4 shows  $d_{\text{norm}}(\rho_{\mathcal{A}^*}, \hat{\rho}_{\mathcal{A}^*})$  over the 20 runs and for varying  $L$ , and the left panel in Figure S2 in the supplementary material displays the proportion of correctly selected items in  $\mathcal{A}^*$  for each  $L$ . The results showed a clear indication that a smaller  $L$  results in a lower error overall. Similarly, the middle panel in Figure 4 (and also the middle panel in Figure S2 in the supplementary material) indicates that the leap size  $l$  has a smaller effect on the error, even if  $l = \text{round}(n^*/5)$  yields a slightly better result, both in terms of the footrule distance ( $d_{\text{norm}}$ ) and in terms of the selection of items. Similar conclusions can be made by inspecting the heatplots of the marginal posterior distributions of  $\rho$  for varying values of the tuning parameters (see Figure S3 in the supplementary material for  $\alpha = 3$ , and Figure S4 in the supplementary material for  $\alpha = 10$ ), where it is evident that varying  $l$  does not affect the error (but does contribute slightly to the degree of uncertainty in the ranks), while varying  $L$  does seem to affect the error to a slight higher degree



**FIGURE 4** Results from the sensitivity study on the tuning parameters described in Section 3.2: Boxplots of  $d_{\text{norm}}$ , on 20 different datasets. From left to right: For varying  $L$ , leap size  $l$ , and  $n_{\text{guess}}^*$ .  $n = 100$ ,  $N = 10$ ,  $n^* = 10$ ,  $\alpha = 5$ ,  $M = 5000$

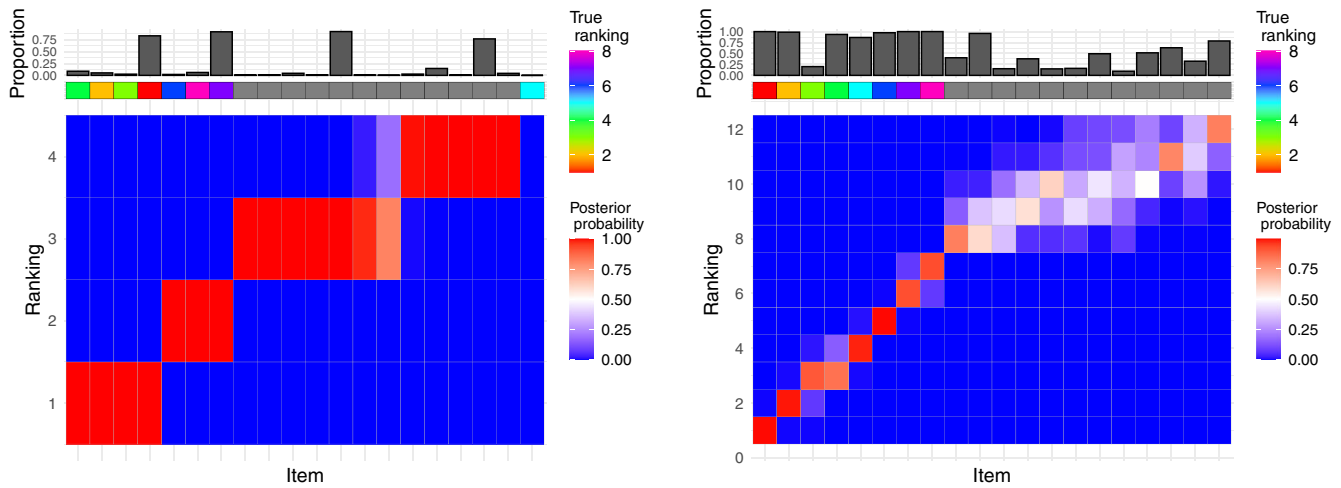


**FIGURE 5** Results from the sensitivity study on the tuning parameters as described in Section 3.2: Acceptance probabilities of the consensus ranking  $\rho$  (left), and of the set of selected items  $\mathcal{A}^*$  (right). The x-axis corresponds to varying  $L = [1, 2, 3]$ , while the points shapes correspond to varying the leap size  $l = [1, 2, 3]$ . Simulated scenario with  $n = 20$ ,  $N = 50$ ,  $n^* = 8$ ,  $\alpha = 3$

compared to  $l$ . Furthermore, as noticeable in the barplots in Figures S3 and S4, increasing  $L$  contributes to worsening the degree of exploration of the items space when estimating the posterior distribution for  $\mathcal{A}^*$ , thus supporting the decision to keep  $L$  small.

It is also worth exploring variations in the acceptance probability of the two parameters estimated in the MCMC,  $\rho$  and  $\mathcal{A}^*$ , when varying the tuning parameters  $L$  and  $l$ . From inspection of Figure 5 ( $\alpha = 3$ ,  $n^* = 8$ ,  $n = 20$ , and  $N = 50$  in this scenario), it seems that the acceptance probability for  $\rho$  increases with increasing  $L$ , while it decreases with increasing  $l$ ; the latter behavior is expected, while the former is less trivial to understand. For what concerns  $\mathcal{A}^*$ , the acceptance probability decreases greatly with increasing  $L$ , and it seems to not be much affected by the leap size  $l$ , which are both quite natural behaviors. Therefore, we generally recommend to keep  $L = 1$  to explore the space appropriately when estimating  $\mathcal{A}^*$ , and to tune  $l$  by inspecting the acceptance probability of  $\rho$ . The tuning of  $l$  in the BMM has been extensively studied, resulting in a recommendation to use  $l = \text{round}(n/5)$  as a default value to ensure a good balance between space exploration and accuracy (see Vitelli et al<sup>17</sup> for further discussions). However, the respective values of  $n^*$  and  $n$ , which constitute the novelty of lowBMM with respect to the original BMM, might make it challenging to choose a combination of tuning parameters that results in a chain that explores the space well. Therefore, even if  $l = \text{round}(n^*/5)$  will generally be a solid choice, we do recommend to try some values of  $l$  before setting it.

Since the number of relevant items  $n^*$  is in general unknown, we also tested the performance of lowBMM when assuming a different  $n^*$  than what was used when generating the data. We did a systematic study with  $n_{\text{guess}}^* \in [6, 8, 10, 12, 20]$ , and with the truth set to  $n^* = 10$ . The results from the study are displayed in the right panels in Figure 4 and in Figure S2 in the supplementary material, both suggesting that setting  $n_{\text{guess}}^* > n^*$  increases the accuracy. This is not surprising, as it allows the algorithm to search for a good selection in a larger permutations space. It is also interesting to evaluate how the consensus ranking  $\rho$  changes when  $n_{\text{guess}}^*$  differs from the true value: the marginal posterior distribution of  $\rho$  obtained when  $n_{\text{guess}}^* < n^*$  is displayed in Figure 6 on the left, and when  $n_{\text{guess}}^* > n^*$  on the right, with  $n^* = 8$ . By inspecting the left panel we see that assuming  $n_{\text{guess}}^* < n^*$  yields a clustered structure, thus implying that the best lower-dimensional solution of ranked data in a higher-dimension is consistent with a partial ordering, also called “bucket solution.”<sup>35</sup> On the other hand, assuming  $n_{\text{guess}}^* > n^*$  (Figure 6, right panel) yields a top-rank solution, as the algorithm seems to be able to



**FIGURE 6** Results from the sensitivity study on the tuning parameters as described in Section 3.2, here when setting  $n_{\text{guess}}^*$  differently from the true  $n^* = 8$ . Both panels display heatplots of the marginal posterior distribution of  $\rho$  with the items ordered according to  $\hat{\rho}_{\mathcal{A}^*}$  on the x-axis; left:  $n_{\text{guess}}^* = 4$ , right:  $n_{\text{guess}}^* = 12$ . The rainbow grid indicates the true ranking  $\rho_{\mathcal{A}^*}$ , and the bar plot indicates the proportion of times the items were selected in  $\mathcal{A}^*$  over all MCMC iterations.  $n = 20$ ,  $N = 50$ ,  $\alpha = 3$ ,  $L = 1$ , and  $l = \text{round}(n_{\text{guess}}^*/5)$

indirectly suggest the true  $n^*$  by showing a much larger uncertainty around the items ranked with larger values than the true  $n^*$ . Overall, when  $n_{\text{guess}}^* > n^*$ , the posterior distribution of  $\rho$  is anyway accurate, since the subset of relevant items has larger marginal posterior probability of being top-ranked. Moreover, the marginal posterior distribution of  $\rho$  implicitly suggests the most suitable value for  $n_{\text{guess}}^*$ , by suddenly showing larger uncertainty around items unnecessarily selected in  $\mathcal{A}^*$ . Similar trends were observed in simulation studies when varying  $\alpha$  and the tuning parameters  $l$  and  $L$ , see Figures S5 and S7 in the supplementary material for  $\alpha = 3$ , and Figures S6 and S8 in the supplementary material for  $\alpha = 10$ , with  $n_{\text{guess}}^* \neq n^*$ .

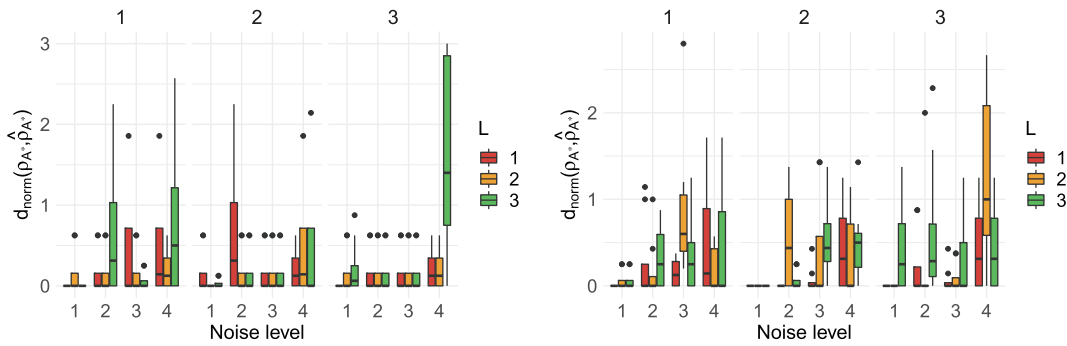
### 3.3 | Robustness to noise experiment

In the simulated data experiments conducted so far, we only tested a varying degree of agreement among the rankings provided by the assessors by varying  $\alpha$ . It could then be argued that both data generating processes were somewhat easy to handle for lowBMM, since the noise structure was only induced by controlling a single model parameter. To test the robustness of lowBMM to increasingly large noise levels, regardless the value of  $\alpha$ , we performed another simulation experiment: data were generated as previously, but an increasingly large random perturbation of the rankings was then enforced in the data. Precisely, we first simulated the data according to the top-rank data generating process described in Section 3.1 (with  $n^* = 8$ ,  $n = 20$ , and  $N = 25$ ), and then we iteratively swapped the rankings of items in the top- $n^*$  (sampled from the Mallows) with those of items ranked below. In order to perform this iterative rank swapping, we followed the following procedure: at level 1, we take the “bottom item” in the true set  $\mathcal{A}^*$  (the item whose rank is  $n^*$  in the true  $\rho_{\mathcal{A}^*}$ ), and we swap the rank for that item with the rank of an item outside of the true set, for 90% of the assessors. Then, at any level  $i > 1$ , we repeat the same procedure for the items in the true set  $\mathcal{A}^*$  whose ranks in  $\rho_{\mathcal{A}^*}$  are  $n^* - i$ . So for instance, if we repeat the swapping procedure four times, at level 4 the ranks of four items in the true set have been swapped with the ranks of four items from the outside. The tuning parameters were distributed over the grid  $(L, l) \in [1, 2, 3] \times [1, 2, 3]$ . The results can be seen in Figure 7, displaying boxplots of  $d_{\text{norm}}(\rho_{\mathcal{A}^*}, \hat{\rho}_{\mathcal{A}^*})$ . The results generally show a very good performance of lowBMM even with an increasing level of noise, that is, several levels in the iterative perturbation/swapping. The figure also suggests that a lower  $L$  is more robust to noise, while the leap size  $l$  has very little effect: therefore, studying lowBMM’s robustness to noise also confirmed the conclusions about the tuning of  $L$  and  $l$  reached in Section 3.2.

### 3.4 | Comparison with other methods

To further assess the accuracy and solidity of the lowBMM method, it is also important to compare its performance with competitor approaches. In order to carry out a fair comparison, we set up simulation experiments with varying






**FIGURE 7** Results from the simulation experiments described in Section 3.3 with  $n = 20$ ,  $N = 25$ ,  $n^* = 8$ : Boxplots of  $d_{\text{norm}}$  for each level of noise added to the data. Left:  $\alpha = 3$ ; right:  $\alpha = 10$ . Each subplot from left to right displays the results obtained when varying the leap size  $l = 1, 2, 3$ , and the color of the bars indicate  $L = 1, 2, 3$

computational demands, taking inspiration from the simulation study in Zhu et al.<sup>20</sup> The first experiment is a toy example with a small number of items,  $n = 20$ , so that all compared approaches can manage to provide a solution in a small setting. Two additional experiments with increasing dimensions, instead, are used to test how the various methods are able to handle a larger number of items, specifically:  $n = 100$  and  $n = 1000$ . The methods included in the comparison were: four Markov chain-based methods  $MC_1$ ,  $MC_2$ ,  $MC_3$ <sup>27</sup> and cross entropy Monte Carlo (CEMC);<sup>28</sup> three Mallows-based methods: the original Mallows model (MM),<sup>16</sup> the extended Mallows model (EMM),<sup>30</sup> and the partition Mallows model (PAMA).<sup>20</sup> Where possible, a comparison with the original BMM<sup>17</sup> was also included. Finally, we included BORDA<sup>36</sup> as a reference, as it is a simple method and quite commonly used as a “baseline” comparison. The TopKLists R package was used for the MC methods, CEMC and BORDA; PerMallows was used for MM; ExtMallows was used for EMM; the PAMA package was used for PAMA; and finally the BayesMallows package was used for BMM. All methods were used to obtain estimates in all simulation scenarios, however some methods did not manage to complete/converge by the set time limit.

We computed three performance measures in order to compare and evaluate the methods. Let  $\rho_{A^*}$  be the true consensus ranking of the relevant items, and  $\hat{\rho}_{A^*}$  be its estimate (the point estimate for the frequentist methods, and the default posterior summary for the Bayesian ones). Similarly, let  $\mathcal{A}^*$  be the true set of relevant items, and  $\hat{\mathcal{A}}^*$  be its estimate, with  $|\hat{\mathcal{A}}^*| = |\mathcal{A}^*| = n^*$ . We then consider the following measures of performance:

1. The footrule distance between the true and the estimated consensus ranking, normalized by the number of correctly selected items  $n_{\text{corr}}$ , denoted as  $d_{\text{norm}}$  (this is the same measure as described in Section 3.2).
2. The coverage, that is, the proportion of correctly selected items  $\hat{p} = n_{\text{corr}}/n^*$ .
3. The recovery distance  $d_R = d_\tau(\rho_{A^*}, \hat{\rho}_{A^*}) + (n^* - n_{\text{corr}}) \times (n + n^* + 1)/2$  as defined in Zhu et al.,<sup>20</sup> where  $d_\tau$  is the Kendall distance between two rankings, and  $n^* - n_{\text{corr}}$  is the number of wrongly selected items.

The results from the first experiment with  $n = 20$  items are shown in the upper portion of Table 1, and in the left panel of Figure 8. BORDA, EMM, and lowBMM all perform similarly, outperforming the other methods in terms of accuracy, both in the selection of the items, and also in their respective rankings. Similar results were obtained for the experiment with  $n = 100$  items, where however we were only able to run the methods displayed in the right panel of Figure 8 in the set time limit of 72 hours (see also the lower portion of Table 1). Finally, concerning the large experiment with  $n = 1000$  items, we simulated data according to the two data generating processes (top-rank and rank consistency) described in Section 3.1. Results are reported in Table 2: only the Mallows-based methods, and not the PAMA, were able to converge in these scenarios. The results for the top-rank simulation experiment are similar to those obtained for the smaller set of items, however with EMM performing slightly better in terms of  $d_R$  and  $\hat{p}$ . On the other hand, EMM was by far the slowest method in this computationally demanding scenario, as the algorithm is based on an iterative optimization procedure and not on MCMC-sampling. In the rank consistency experiment, lowBMM is performing slightly better than EMM for  $d_R$  and  $\hat{p}$  (see Figure 9, and the right portion of Table 2). These results also show that BMM and MM are not even close to converging to the correct solution in the rank consistency scenario. It is also worth noting that lowBMM and BMM are Bayesian methods that provide uncertainty quantification, while EMM and MM are methods that result in point estimates.



**FIGURE 8** Results of the top-rank simulation experiments described in Section 3.4: Boxplots of the values of  $d_{\text{norm}}$  over 50 repetitions obtained with several alternative methods. Left panel:  $\alpha = 2$ ,  $n = 20$ ,  $N = 5$ ,  $n^* = 8$ ,  $M = 1000$ ; right panel:  $\alpha = 2$ ,  $n = 100$ ,  $N = 10$ ,  $n^* = 10$ ,  $M = 2000$

**TABLE 1** Average results over 50 repetitions of the top-rank simulation experiments described in Section 3.4

	BORDA	CEMC	EMM	lowBMM	MC1	MC2	MC3	MM	PAMA
<i>Top-rank, <math>n = 20</math></i>									
$d_R$	<b>12.24</b>	80.71	<b>11.02</b>	<b>12.06</b>	72.52	42.78	27.26	80.53	38.14
$\hat{p}$	<b>1.00</b>	0.43	<b>1.00</b>	<b>1.00</b>	0.50	0.75	0.88	0.42	0.78
$d_{\text{norm}}$	5.29	19.03	<b>4.87</b>	5.17	14.07	8.35	6.64	18.60	7.87
Time (s)	0.02	58.42	0.05	0.54	4.42	4.42	4.42	0.0003	17.58
<i>Top-rank, <math>n = 100</math></i>									
$d_R$	<b>19.88</b>	-	<b>19.40</b>	21.70	-	-	-	525.35	249.33
$\hat{p}$	<b>1.00</b>	-	<b>1.00</b>	<b>1.00</b>	-	-	-	0.09	0.59
$d_{\text{norm}}$	<b>25.79</b>	-	<b>24.97</b>	<b>25.10</b>	-	-	-	Inf	93.28
Time (s)	0.64	-	3.19	1.48	-	-	-	0.01	876.17

Note: Upper portion of the table:  $N = 5$ ,  $n = 20$ ,  $n^* = 8$ ,  $\alpha = 2$ ,  $M = 1000$ ; lower portion of the table:  $N = 10$ ,  $n = 100$ ,  $n^* = 10$ ,  $\alpha = 2$ ,  $M = 2000$ . Bold values represent the best results among the methods in the comparison.


**TABLE 2** Average results over 50 repetitions of the large simulation experiment described in Section 3.4 with  $N = 50$ ,  $n = 1000$ ,  $n^* = 50$ ,  $M = 7.5 \times 10^4$  for the two data generating processes: A top-rank simulation experiment with  $\alpha = 2$  (left portion of the table), and a rank-consistency simulation experiment with  $\alpha = 5$  (right portion of the table)

	Top-rank				Rank consistency			
	BMM	EMM	LowBMM	MM	BMM	EMM	LowBMM	MM
$d_R$	9269.69	<b>611.10</b>	1621.02	25 550.13	23 344.03	14 776.84	<b>12 933.85</b>	25 518.04
$\hat{p}$	0.67	<b>1.00</b>	0.96	0.05	0.13	0.45	<b>0.53</b>	0.05
$d_{\text{norm}}$	594.76	<b>323.83</b>	336.73	Inf	3105.60	<b>398.23</b>	557.30	Inf
Time (s)	116.36	2307.42	236.92	9.48	105.41	1586.66	227.79	8.76

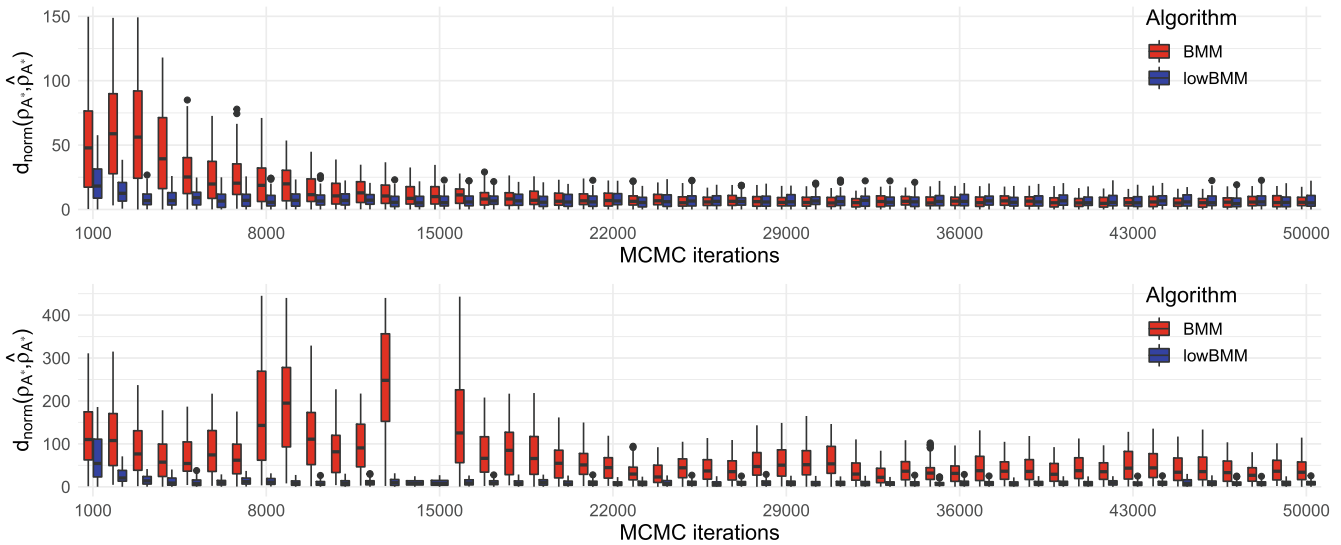
Note: Bold values represent the best results among the methods in the comparison.

### 3.4.1 | Comparison with BMM

Among the performance comparisons of lowBMM with alternative methods, a special role is played by BMM as it is the original full dimensional version of lowBMM. Thus, we would also like to verify that lowBMM is consistent with the full model in a limiting situation. To this aim, we ran the full BMM as well as the variable selection lowBMM on the same simulated dataset with  $n = 20$ ,  $N = 50$ ,  $n^* = 8$ . For a comparison of the marginal posterior distribution of  $\rho$  obtained on the reduced set of items with the lowBMM to the one obtained with BMM on the complete set of items, see Figure S10



**FIGURE 9** Results of the rank consistency simulation experiment described in Section 3.4 with  $n = 1000, N = 50, n^* = 50, M = 7.5 \times 10^4, \alpha = 5$ : Boxplots of the values of  $\hat{p}$  (left) and  $d_{\text{norm}}$  (right) over 50 repetitions obtained with several alternative methods (BMM, EMM, lowBMM, and MM)



**FIGURE 10** Results from the simulation experiment described in Section 3.4.1 with  $n = 1000, N = 50, n^* = 50$ : Boxplots of the posterior distribution of  $d_{\text{norm}}$  evolving along MCMC iterations for BMM (red) and lowBMM (blue). Top panel: Top-rank simulation scenario with  $\alpha = 5$ ; bottom panel: rank-consistency simulation scenario with  $\alpha = 2$

in the supplementary material, right and left panel, respectively. Both models show consistent results for the top items, with the full model being slightly more certain in its top ranks. The noisy items seem to be randomly distributed on the bottom in the full model, while in lowBMM they are assigned the bottom ranks.

We then compared BMM and lowBMM on a larger simulated scenario ( $n = 1000, N = 50, n^* = 50$ ), when using both data generating processes: Figure 10 demonstrates that lowBMM outperforms BMM in terms of accuracy and stability in both scenarios. In the top-rank simulation experiment, lowBMM converges to the correct solution with fewer iterations, while BMM reaches the solution at a later stage. This shows the consistency of lowBMM, as we obtain the same solution as for the full model, however in the correct reduced dimension and with lower computational demands. For the rank-consistency simulation experiment, lowBMM clearly outperforms BMM: while lowBMM converges to the solution after 3000 iterations, BMM exhibits an oscillating behavior, and it never manages to reach convergence. This last simulation example further shows the flexibility in the variable selection procedure under different data generating procedures, even in a high-dimensional setting with basically no possible (Bayesian) competitor.

### 3.5 | Off-line estimation of $\alpha$

Finally, we also aimed at testing the approach described in Section 2.4 for the off-line estimation of  $\alpha$ , in a simulation setting with the same top-rank data generating process as previously used. We therefore generated datasets with varying

**TABLE 3** Results from simulation experiments described in Section 3.5, for testing the off-line estimation of  $\alpha$  for varying  $(N, n)$  as described in Section 2.4

$N$		10	20	50
$n$	$n^*$	$\hat{\alpha}_{n^*}$		
20	7	2.37	2.22	2.18
150	50	2.15	2.18	2.23
500	167	2.14	2.15	2.24
1000	333	2.05	2.09	2.21

Note:  $\alpha_{\text{true}} = 3$  for all scenarios.

dimensions  $(N, n)$ , where  $n^* = \text{round}(n/3)$  items were sampled from the Mallows model with a fixed  $\alpha_{\text{true}}$ , and the rest of the items were sampled from a uniform distribution. For each dataset in a given dimension, we ran the estimation procedure of Section 2.4 by generating datasets of equal dimension over a grid of  $\alpha_0$  values, and computed the mean distance between assessors in each dataset. Finally, we obtained the final estimate  $\hat{\alpha}_{n^*}$  by determining the intersection between the distance measure for the simulated datasets and the distance measure for the dataset generated from  $\alpha_{\text{true}}$ , in each dimension. Performance of the estimation procedure can be inspected in Table 3, where the method stability is demonstrated even for increasing  $n$ . The estimated values  $\hat{\alpha}_{n^*}$  are consistently slightly smaller than  $\alpha_{\text{true}}$ , which however does not affect the method performance, but simply makes the posterior distributions for  $\rho$  and  $\mathcal{A}^*$  slightly more peaked to absorb the unnecessary uncertainty. We remark that a smaller  $\alpha$  than necessary is in theory better for estimation, as it allows for more model flexibility and for better space exploration in the MCMC.

## 4 | APPLICATION TO RNAseq DATA FROM OVARIAN CANCER PATIENTS


In this section, we describe the application of the lowBMM method on RNAseq data from TCGA ovarian cancer patients. The gene level expression data was available at the TCGA Data Portal (<https://tcga-data.nci.nih.gov/tcga/>),<sup>‡</sup> and the selected ovarian cancer patients were the same as previously analyzed.<sup>2</sup>

Serous ovarian tumors share similar genetic origin and features with basal-like breast cancers:<sup>39</sup> both are difficult-to-treat cancer types, and they are highly commonly characterized at the molecular level, especially for what concerns mutations (with TP53 being the most frequent example). Massive whole-genome data analyses are therefore key to understanding individual tumor biology for these cancer types. This is the main motivation for focusing on RNAseq data from serous ovarian tumors, as selecting basal-like breast cancer patients would require an additional subtyping preprocessing step that we would like to avoid. Additionally, since our method assumes complete data, preprocessing steps included keeping only the items (genes) with less than 50% missing values: all of these missing values were imputed using k-nearest neighbor averaging.<sup>40</sup> This resulted in a final dataset of ultra-high dimension with  $n = 15\,348$  genes and  $N = 265$  patients.


For running our lowBMM method, we set  $n^* = 500$ ,  $l = 20$ , and  $L = 1$  (tuning parameters were set based on the conclusions from the simulation studies in Section 3.2, and from some trial-and-error). We also performed an off-line estimation of  $\alpha$  on the experimental dataset, on a search grid  $\alpha \in [10^{-15}, 10^{-5}, 1, 10, 50, 100]$  that was set as such to ensure adequate coverage over several orders of magnitude, as described in Section 2.4, resulting in  $\alpha = 10$ . Due to the high-dimensionality of the data, we ran two chains, each with  $M = 5 \times 10^6$  MCMC iterations on a big memory HPC server. We discarded the first  $5 \times 10^4$  iterations as burn-in in both chains before merging them for post-processing (see Section 2.3.1 for details on the post-processing of the results). As  $\mathcal{A}^*$  is a probabilistic selection, we will get slightly different posterior summaries from different runs of lowBMM. Moreover,  $\mathcal{A}^*$  is a discrete parameter defined in a space whose dimension grows exponentially with  $n$ , and in this particular case the space will be exceedingly large. This will in turn result in small variations in each run of the algorithm, as we are limited by computational capacities in how much we can realistically explore this space. Furthermore, this behavior can also be affected by aspects related to the current problem setting: particularly as the  $n^*$  value used in the analysis might be too low compared to  $n$ . We remark that the sensitivity study on the tuning parameters in Section 3.2 suggested that running lowBMM with a smaller  $n^*$  compared to the “truth” might make it more

difficult for the method to converge to the true  $\rho_{\mathcal{A}^*}$ . Moreover, during the off-line estimation of  $\alpha$ , we estimated the  $n^*$  that best fits the experimental data to  $n^* \approx 2400$ . However, we kept the final  $n^* = 500$  value as it is more suited to allow results interpretability, and as it also allows the MCMC to perform a good exploration of the state space given the current computational capacity. The computations were performed on resources provided by UNINETT Sigma2, the National Infrastructure for High-Performance Computing and Data Storage in Norway.

To gain insight on the results provided by lowBMM on the experimental dataset, we computed a “top probability selection” as described in Section 2.3.1, to be used in a Gene Set Enrichment Analysis (GSEA).<sup>41,42</sup> This analysis was done to verify whether the method could identify common genes related to cancer. To this aim, we computed various



**FIGURE 11** Results from the analysis of the RNaseq data described in Section 4. Left: Violin plots of  $P(A_i \in \text{top-}K \mid A_i \in \hat{\mathcal{A}}^*)$  for all items  $A_i, i = 1, \dots, n$  and for varying  $K$ . Right: Histogram of  $P(A_i \in \text{top-}75 \mid A_i \in \hat{\mathcal{A}}^*)$  for all items  $A_i, i = 1, \dots, n$ , with dashed line at  $c = 0.77$ .  $n^* = 500, \alpha = 10, L = 1, l = 10, M = 10^7$ , and burn-in =  $10^5$



**FIGURE 12** Results from the analysis of the RNaseq data described in Section 4: Heatplot of the marginal posterior of  $\rho$  depicted only for the items in the “top probability selection”  $\hat{\mathcal{A}}_{\text{top}}^*$ , and with the items ordered as in  $\hat{\rho}_{\mathcal{A}^*}$  on the x-axis. The bars on top indicate the proportion of times the items were selected in  $\mathcal{A}^*$  over all MCMC iterations, with  $n^* = 500, \alpha = 10, L = 1, l = 10, M = 10^7$ , and burn-in =  $10^5$



**TABLE 4** Overview of the top-10 gene sets ranked according to the associated  $P$ -value from a GSEA performed on the selection  $\hat{A}_{\text{top}}^*$ ,  $|\hat{A}_{\text{top}}^*| = 63$

Gene set name	# genes in gene set	# genes in overlap	$P$ -value
Regulation of cell differentiation (GOBP)	1618	16	2.59E-9
Regulation of multicellular organismal development (GOBP)	1397	13	2.05E-7
Regulation of anatomical structure morphogenesis (GOBP)	1006	11	4.17E-7
Positive regulation of developmental process (GOBP)	1284	12	6.17E-7
Positive regulation of cell differentiation (GOBP)	844	10	7.2E-7
Response to endogenous stimulus (GOBP)	1624	13	1.12E-6
Cellular response to nitrogen compound (GOBP)	698	9	1.37E-6
Sensory organ development (GOBP)	534	8	1.84E-6
Animal organ morphogenesis (GOBP)	1025	10	4.08E-6
Striated muscle cell differentiation (GOBP)	269	6	4.12E-6

top- $K$  sets, with a threshold for the lower items: genes that had probability less than 0.1 of being in the different top- $K$  sets were dropped from the plot. This resulted in the violin plots shown in Figure 11, left panel: this plot indicates that  $K = 75$  is a good choice for allowing many genes to have a quite large posterior probability of being top-ranked. We therefore further inspected the probability distribution of the top-75 probabilities across the genes, by drawing a histogram of the same probabilities (Figure 11, right panel): this second plot allowed us to select a cut-off point of  $c = 0.77$  for the top-75 probability set, thus resulting in a final selection of  $|\hat{A}_{\text{top}}^*| = 63$  genes where  $\hat{A}_{\text{top}}^* = \{A_i \in \mathcal{A}^* \text{ s.t. } P(A_i \in \text{top-75}, i = 1, \dots, n \mid A_i \in \hat{A}^*) > 0.77\}$ .

Figure 12 shows that the items in the refined selection  $\hat{A}_{\text{top}}^*$  are selected with very large posterior probability, even if they sometimes compete for some of the available ranks as is often expected in ultra-high dimensions. This behavior might also be due to  $n^*$  being smaller than  $n_{\text{true}}^*$ , as the choice of  $n^*$  will be constrained by computational capacities in such dimensions. There is overall more variability in the rankings of the items as compared to the simulation studies, which is to some degree expected in experimental data with a higher degree of noise. Higher variability in the ranks was also typical for the simulation studies where  $n_{\text{guess}}^* \ll n_{\text{true}}^*$  (see Figure S5 in the supplementary material).

The GSEA indicated that the lowBMM identified known genes involved in cell differentiation and cancer development. The top-ranked genes selected in  $\hat{A}_{\text{top}}^*$  were enriched in pathways such as cell differentiation, morphogenesis, and developmental processes (Table 4), suggesting a role in the deregulation of cellular identity essential for cancer development.

## 5 | DISCUSSION

In this article, we have developed a novel rank-based unsupervised variable selection method for high-dimensional data. The method extends the applicability of the Bayesian Mallows ranking model to unprecedented data dimensions and provides a much better modeling approach. Indeed, an important advantage of the variable selection procedure is the ability to distinguish relevant genes from the background ones, and to provide an estimate of the relative importance of the selected relevant genes. Furthermore, the variable selection procedure is able to work in high-dimensional settings where the number of items is  $n \gg 10^3$ , with no competitor model for ranking data capable of scaling to such setting. The simulation studies described in Section 3 showed that the method performs well on datasets of varying sizes and under varying data generating procedures. Additionally, lowBMM is superior in terms of accuracy and computational time compared to existing methods in high dimensions, and has no competitor in *ultra*-high dimensions. We applied the proposed model to RNAseq data of ovarian cancer samples from TCGA (Section 4). Although we consider a cancer genomics application, the methodological contribution given in this article is very general, and can thus be applied to any high-dimensional setting with heterogeneous rank data.

Throughout this work, we have assumed that the number of relevant items to be selected,  $n^*$ , is known. This is reasonable in many practical problems where a specific  $n^*$  can be tuned according to the research

questions/demands or has to be decided according to computational capacity. Simulation studies showed that assuming  $n_{\text{guess}}^* \ll n_{\text{true}}^*$  might pose problems in high-dimensional settings, and therefore the joint estimation of this parameter as part of the existing hierarchical framework constitutes a very interesting direction for future research.

Nonetheless, as mentioned in the simulation studies while commenting upon the marginal posterior distribution of  $\rho$  obtained when  $n_{\text{guess}}^* < n_{\text{true}}^*$ , our method can be seen as an alternative to the constrained median bucket order technique.<sup>35</sup> It would then be interesting to fully compare the respective estimates obtained via lowBMM and the median bucket order, in the light of the relationship between the choice of permutation space (strong vs weak rankings), distance (any distance vs the Kemeny only), and estimation procedure (Bayesian vs frequentist). This comparison is clearly out of the scopes of the present article, but makes up a very interesting future research direction.

The current version of lowBMM works under several assumptions: no handling of missing data, no clustering, fixed scale parameter  $\alpha$ , all aspects that we plan to implement in future versions of the method. The somewhat challenging procedure of choosing the tuning parameters  $l$  and  $L$  might be solved with an adaptive MCMC and would allow the algorithm to explore the space more thoroughly. We leave this task for future consideration.

To our knowledge, no rank-based unsupervised variable selection procedure that is capable of scaling to the common data dimensions in -omics applications currently exists.

## AUTHOR CONTRIBUTIONS

Emilie Eliseussen and Valeria Vitelli defined the theoretical and computational aspects of the method together and drafted the article. Thomas Fleischer provided support with the GSEA and helped with the interpretation of the results.

## ACKNOWLEDGEMENTS

We thank Prof. Arnoldo Frigessi for fruitful discussions and Dr. Manuela Zucknick for help with preprocessing of the RNAseq data.

## CONFLICT OF INTEREST

The authors declare no potential conflict of interest.

## ENDNOTES

\*This feature is unobserved in most recommender system applications of the model, where we assume the users provide their ranking/preference according to the feature levels in their mind. On the other hand, when continuous data are converted to ranks for better modeling (eg, in transcriptomic data analyses), the feature corresponds to the observed continuous measurement.

†A right-invariant metric  $d(\cdot, \cdot)$  is such that, for any  $r_1, r_2 \in \mathcal{P}_n$ , it holds:  $d(r_1, r_2) = d(r_1 r_2^{-1}, \mathbf{1}_n)$ ,  $\mathbf{1}_n = \{1, 2, \dots, n\}$ .<sup>31</sup>

‡The BAM files were downloaded from CGHub ([www.cghub.ucsc.edu](http://www.cghub.ucsc.edu)), and converted to FASTQs, where all samples were processed similarly aligning reads to the hg19 genome assembly using MapSplice.<sup>37</sup> Gene expression was quantified for the transcript models corresponding to the TCGA GAF 2.13, using RSEM4 and normalized within-sample to a fixed upper quartile.<sup>38</sup>

## ORCID

Emilie Eliseussen  <https://orcid.org/0000-0003-0465-7221>

## REFERENCES

- Alexandrov L, Nik-Zainal S, Wedge D, Campbell P, Stratton M. Deciphering signatures of mutational processes operative in human cancer. *Cell Rep*. 2013;3(1):246-259. doi:10.1016/j.celrep.2012.12.008
- Hoadley K, Yau C, Wolf D, et al. Multiplatform analysis of 12 cancer types reveals molecular classification within and across tissues of origin. *Cell*. 2014;158(4):929-944. doi:10.1016/j.cell.2014.06.049
- Hoadley K, Yau C, Hinoue T, et al. Cell-of-origin patterns dominate the molecular classification of 10,000 tumors from 33 types of cancer. *Cell*. 2018;173(2):291-304.e6. doi:10.1016/j.cell.2018.03.022
- Ickstadt K, Schäfer M, Zucknick M. Toward integrative Bayesian analysis in molecular biology. *Annu Rev Stat Appl*. 2018;5(1):141-167. doi:10.1146/annurev-statistics-031017-100438
- Bühlmann P, van de Geer S. *Statistics for High-Dimensional Data*. Springer Series in Statistics. Heidelberg, Germany: Springer; 2011.
- Ideker T, Dutkowski J, Hood L. Boosting signal-to-noise in complex biology: prior knowledge is power. *Cell*. 2011;144(6):860-863. doi:10.1016/j.cell.2011.03.007

7. Bouveyron C, Latouche P, Mattei PA. Bayesian variable selection for globally sparse probabilistic PCA. *Electron J Stat.* 2018;12(2):3036-3070. doi:10.1214/18-EJS1450
8. Agrawal A, Chiu AM, Le M, Halperin E, Sankararaman S. Scalable probabilistic PCA for large-scale genetic variation data. *PLoS Genet.* 2020;16(5):1-19. doi:10.1371/journal.pgen.1008773
9. Li F, Zhang NR. Bayesian variable selection in structured high-dimensional covariate spaces with applications in genomics. *J Am Stat Assoc.* 2010;105(491):1202-1214. doi:10.1198/jasa.2010.tm08177
10. Kundu S, Cheng Y, Shin M, Manyam G, Mallick BK, Baladandayuthapani V. Bayesian variable selection with graphical structure learning: applications in integrative genomics. *PLoS One.* 2018;13(7):1-29. doi:10.1371/journal.pone.0195070
11. Deng K, Han S, Li K, Liu J. Bayesian aggregation of order-based rank data. *J Am Stat Assoc.* 2014;109(507):1023-1039. doi:10.1080/01621459.2013.878660
12. Afsari B, Braga-Neto U, Geman D. Rank discriminants for predicting phenotypes from RNA expression. *Ann Appl Stat.* 2014;8(3):1469-1491. doi:10.1214/14-AOAS738
13. Alvo M, Yu PL. *Statistical Methods for Ranking Data.* Frontiers in Probability and the Statistical Sciences. New York, NY: Springer; 2014.
14. Luce RD. *Individual Choice Behavior: A Theoretical Analysis.* New York, NY: Wiley; 1959.
15. Plackett RL. The analysis of permutations. *J R Stat Soc Ser C Appl Stat.* 1975;24(2):193-202. doi:10.2307/2346567
16. Mallows CL. Non-null ranking models. *Biometrika.* 1957;44(1/2):114-130. doi:10.1093/biomet/44.1-2.114
17. Vitelli V, Sørensen O, Crispino M, Frigessi A, Arjas E. Probabilistic preference learning with the mallows rank model. *J Mach Learn Res.* 2017;18(158):1-49.
18. Sørensen O, Crispino M, Liu Q, Vitelli V. BayesMallows: an R package for the Bayesian mallows model. *R J.* 2020;12(1):324-342. doi:10.32614/RJ-2020-026
19. Liu Q, Crispino M, Scheel I, Vitelli V, Frigessi A. Model-based learning from preference data. *Annu Rev Stat Appl.* 2019;6(1):329-354. doi:10.1146/annurev-statistics-031017-100213
20. Zhu W, Jiang Y, Liu JS, Deng K. Partition-mallows model and its inference for rank aggregation. *J Am Stat Assoc.* 2021;1-41. doi:10.1080/01621459.2021.1930547
21. Cui C, Singh SP, Staicu AM, Reich BJ. Bayesian variable selection for high-dimensional rank data. *Environmetrics.* 2021;32(7):1-16. doi:10.1002/env.2682
22. Koop G, Poirier DJ. Rank-ordered logit models: an empirical analysis of ontario voter preferences. *J Appl Econ.* 1994;9(4):369-388.
23. Yao G, Bockenholt U. Bayesian estimation of Thurstonian ranking models based on the Gibbs sampler. *Br J Math Stat Psychol.* 1999;52(1):79-92. doi:10.1348/000711099158973
24. Yu P. Bayesian analysis of order-statistics models for ranking data. *Psychometrika.* 2000;65(3):281-299. doi:10.1007/BF02296147
25. Johnson V, Deaner R, Schaik C. Bayesian analysis of rank data with application to primate intelligence experiments. *J Am Stat Assoc.* 2002;97(457):8-17. doi:10.2307/3085754
26. Barney B, Amici F, Aureli F, Call J, Johnson V. Joint Bayesian modeling of binomial and rank data for primate cognition. *J Am Stat Assoc.* 2015;110(510):573-582. doi:10.1080/01621459.2015.1016223
27. Lin S. Space oriented rank-based data integration. *Stat Appl Genet Mol Biol.* 2010;9(1):Article20. doi:10.2202/1544-6115.1534
28. Lin S, Ding J. Integration of ranked lists via cross entropy Monte Carlo with applications to mRNA and microRNA Studies. *Biometrics.* 2009;65(1):9-18.
29. Irurozki E, Calvo B, Lozano JA. Mallows and generalized Mallows model for matchings. *Bernoulli.* 2019;25(2):1160-1188. doi:10.3150/17-BEJ1017
30. Li H, Xu M, Liu JS, Fan X. An extended mallows model for ranked data aggregation. *J Am Stat Assoc.* 2020;115(530):730-746. doi:10.1080/01621459.2019.1573733
31. Diaconis P. *Group Representations in Probability and Statistics.* Lecture Notes - Monograph Series. Vol 11. Hayward, CA: Institute of Mathematical Statistics; 1988.
32. Ishwaran H, James LF. Gibbs sampling methods for stick-breaking priors. *J Am Stat Assoc.* 2001;96(453):161-173.
33. Ramanan VK, Shen L, Moore JH, Saykin AJ. Pathway analysis of genomic data: concepts, methods, and prospects for future development. *Trends Genet.* 2012;28(7):323-332.
34. Liu Q, Vitelli V, Mannino C, Frigessi A, Scheel I. Pseudo-Mallows for Efficient Preference Learning arXiv, 2022. doi: 10.48550/ARXIV.2205.13911
35. D'Ambrosio A, Iorio C, Staiano M, Siciliano R. Median constrained bucket order rank aggregation. *Comput Stat.* 2019;34(2):787-802. doi:10.1007/s00180-018-0858-z
36. Borda dJC. Mémoire sur les élections au scrutin. *Histoire de l'Académie Royale des Sciences* 1781.
37. Wang K, Singh D, Zeng Z, et al. MapSplice: accurate mapping of RNA-seq reads for splice junction discovery. *Nucl Acids Res.* 2010;38(18):e178. doi:10.1093/nar/gkq622
38. Li B, Dewey C. RSEM: accurate transcript quantification from RNA-Seq data with or without a reference genome. *BMC Bioinform.* 2011;12:323. doi:10.1186/1471-2105-12-323
39. Begg CB, Rice MS, Zabor EC, Tworoger SS. Examining the common aetiology of serous ovarian cancers and basal-like breast cancers using double primaries. *Br J Cancer.* 2017;116(8):1088-1091. doi:10.1038/bjc.2017.73
40. Troyanskaya O, Cantor M, Sherlock G, et al. Missing value estimation methods for DNA microarrays. *Bioinformatics.* 2001;17(6):520-525. doi:10.1093/bioinformatics/17.6.520

41. Subramanian A, Tamayo P, Mootha VK, et al. Gene set enrichment analysis: a knowledge-based approach for interpreting genome-wide expression profiles. *Proc Natl Acad Sci*. 2005;102(43):15545-15550. doi:[10.1073/pnas.0506580102](https://doi.org/10.1073/pnas.0506580102)
42. Mootha V, Lindgren C, Eriksson KF, et al. PGC-1 alpha-responsive genes involved in oxidative phosphorylation are coordinately downregulated in human diabetes. *Nat Genet*. 2003;34:267-273. doi:[10.1038/ng1180](https://doi.org/10.1038/ng1180)

## SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

**How to cite this article:** Eliseussen E, Fleischer T, Vitelli V. Rank-based Bayesian variable selection for genome-wide transcriptomic analyses. *Statistics in Medicine*. 2022;41(23):4532-4553. doi: [10.1002/sim.9524](https://doi.org/10.1002/sim.9524)