Constrained Extended Plackett-Luce model for the analysis of preference rankings

Cristina Mollica and Luca Tardella

Abstract Choice behavior and preferences typically involve numerous and subjective aspects that are difficult to be identified and quantified. For this reason, their exploration is frequently conducted through the collection of ordinal evidence in the form of ranking data. Multistage ranking models, including the popular Plackett-Luce distribution (PL), rely on the assumption that the ranking process is performed sequentially, by assigning the positions from the top to the bottom one (*forward order*). A recent contribution to the ranking literature relaxed this assumption with the addition of the discrete *reference order* parameter, yielding the novel *Extended Plackett-Luce model* (EPL). In this work, we introduce the EPL with order constraints on the reference order parameter and a novel diagnostic tool to assess the adequacy of the EPL parametric specification. The usefulness of the proposal is illustrated with an application to a real dataset.

Key words: Ranking data, Plackett-Luce model, Bayesian inference, Data augmentation, Gibbs sampling, Metropolis-Hastings, model diagnostics

1 Introduction

A *ranking* $\pi = (\pi(1),\dots,\pi(K))$ of K items is a sequence where the entry $\pi(i)$ indicates the rank attributed to the i-th alternative. Data can be equivalently collected in the ordering format $\pi^{-1} = (\pi^{-1}(1),\dots,\pi^{-1}(K))$, such that the generic component $\pi^{-1}(j)$ denotes the item ranked in the j-th position. Regardless of the adopted

Cristina Mollica
Dipartimento di Metodi e Modelli per il Territorio, l'Economia e la Finanza, Sapienza Università di Roma e-mail: `cristina.mollica@uniroma1.it`

Luca Tardella
Dipartimento di Scienze Statistiche, Sapienza Università di Roma e-mail: `luca.tardella@uniroma1.it`

format, ranked observations are multivariate and, specifically, correspond to permutations of the first K integers.

The statistical literature concerning ranked data modeling and analysis is reviewed in [3] and, more recently, in [1]. Several parametric distributions on the set of permutations $\mathscr{S}_K$ have been developed and applied to real experiments. A popular parametric family is the *Plackett-Luce* model (PL), belonging to the class of the so-called *stagewise ranking models*. The basic idea is the decomposition of the ranking process into $K-1$ stages, concerning the attribution of each position according to the *forward order*, that is, the ordering of the alternatives proceeds sequentially from the most-liked to the least-liked item. The implicit forward order assumption has been relaxed by [4] in the *Extended Plackett-Luce model* (EPL). The PL extension relies on the introduction of the *reference order* parameter indicating the rank assignment order. In this work, we investigate a restricted version of the EPL with order constraints for the reference order parameter representing a meaningful rank attribution process and we also introduce a novel diagnostic to assess the adequacy of the EPL assumption as the actual sampling distribution of the observed rankings.

2 The Extended Plackett-Luce model with order constraints

2.1 Model specification

The implicit assumption in the PL scheme is the forward ranking order, meaning that at the first stage the ranker reveals the item in the first position (most-liked alternative), at the second stage she assigns the second position and so on up to the last rank (least-liked alternative). [4] suggested the extension of the PL by relaxing the canonical forward order assumption, in order to explore alternative meaningful ranking orders for the choice process and to increase the flexibility of the PL parametric family. Their proposal was realized by representing the ranking order with an additional model parameter $\rho = (\rho(1), \ldots, \rho(K))$, called reference order, where the entry $\rho(t)$ indicates the rank attributed at the t-th stage of the ranking process. Thus, ρ is a discrete parameter given by a permutation of the first K integers and the composition $\eta^{-1} = \pi^{-1}\rho$ of an ordering with a reference order yields the sequence $\eta^{-1} = (\eta^{-1}(1), \ldots, \eta^{-1}(K))$ which lists the items in order of selection, such that the component $\eta^{-1}(t) = \pi^{-1}(\rho(t))$ corresponds to the item chosen at stage t and receiving rank $\rho(t)$. The probability of a generic ordering under EPL can be written as

$$\mathbf{P}_{\text{EPL}}(\pi^{-1}|\rho,\underline{p}) = \mathbf{P}_{\text{PL}}(\pi^{-1}\rho|\underline{p}) = \prod_{t=1}^{K} \frac{p_{\pi^{-1}(\rho(t))}}{\sum_{v=t}^{K} p_{\pi^{-1}(\rho(v))}} \qquad \pi^{-1} \in \mathscr{S}_K, \tag{1}$$

Hereinafter, we will shortly refer to (1) as $\text{EPL}(\rho,\underline{p})$. The quantities p_i's are the support parameters and are proportional to the probabilities for each item to be ranked in the position indicated by the first entry of ρ.

Differently from [4], we focus on a restriction $\tilde{\mathscr{S}}_K$ of the whole permutation space S_K for the reference order parameter. Our choice can be explained by the fact that, in a preference elicitation process, not all the possible $K!$ orders seem to be equally natural, hence plausible. Often the ranker has a clearer perception about her extreme preferences (most-liked and least-liked items), rather than middle positions. In this perspective, the rank attribution process can be regarded as the result of a sequential "top-or-bottom" selection of the positions. At each stage, the ranker specifies either her best or worst choice among the available positions at that given step. With this scheme, the reference order can be equivalently represented as a binary sequence $\underline{W} = (W_1, \dots, W_K)$ where the generic W_t component indicates whether the ranker makes a top or bottom decision at the t-th stage, with the convention that $W_K = 1$. One can then formalize the mapping from the restricted permutation ρ to $\underline{W}$ with the help of a vector of non negative integers $\underline{F} = (F_1, \dots, F_K)$, where F_t represents the number of top positions assigned before stage t. In fact, by starting from positing by construction $F_1 = 0$, one can derive sequentially

$$W_t = I_{[\rho(t)=\rho_{\mathrm{F}}(F_t+1)]} = \begin{cases} 1 & \text{at stage } t \text{ the top preference is specified,} \\ 0 & \text{at stage } t \text{ the bottom preference is specified,} \end{cases}$$

where $I_{[E]}$ is the indicator function of the event E and $F_t = \sum_{\nu=1}^{t-1} W_\nu$ for $t = 2, ..., K$. Note that, since the forward and backward orders $(\rho_{\mathrm{F}}, \rho_{\mathrm{B}})$ can be regarded as the two extreme benchmarks in the sequential construction of ρ, this allows us to understand that $\rho_F(F_t+1)$ corresponds to the top position available at stage t. Conversely, $B_t = (t-1) - F_t$ is the number of bottom positions assigned before stage t and thus, symmetrically, one can understand that $\rho_B(B_t+1)$ indicates the bottom position available at stage t.

The binary representation of the reference order suggests that, under the constraints of the "top-or-bottom" scheme, the size of $\tilde{\mathscr{S}}_K$ is equal to 2^{K-1}. The reduction of the reference order space into a finite set with an exponential size, rather than with a factorial cardinality, is convenient for at least two reasons: i) it leads to a more intuitive interpretation of the support parameters, since they become proportional to the probability for each item to be ranked either in the first or in the last position and ii) it facilitates the construction of a Metropolis-Hastings (MH) step to sample the reference order parameter.

2.2 *Bayesian estimation of the EPL via MCMC*

Inference on the EPL and its generalization into a finite mixture framework was originally addressed from the frequentist perspective in [4]. Here we consider the original MCMC methods recently developed by [6] to solve Bayesian inference for the constrained EPL.

In the Bayesian domain, the data augmentation with the latent quantitative variables $\underline{y} = (y_{st})$ for $s = 1, \dots, N$ and $t = 1, \dots, K$ crucially contributes to make it tractable analytically the inference for the EPL . The auxiliary variables y_{st}'s are assumed to be conditionally independent on each other and exponentially distributed with rate parameter equal to the normalization term of the EPL, see also [5]. For the prior specification, independence of $\underline{p}$ and ρ is assumed together with independent Gamma densities for the support parameters, motivated by the conjugacy with the model, and a discrete uniform distribution on $\tilde{\mathscr{S}}_K$ for the reference order. [6] presented a tuned joint Metropolis-within-Gibbs sampling (TJM-within-GS) to perform approximate posterior inference, where the simulation of the reference order is accomplished with a MH algorithm relying on a joint proposal distribution on ρ and $\underline{p}$, whereas the posterior drawings of the latent variables y's and the support parameters are performed from the related full-conditional distributions. At the generic iteration $l+1$, the TJM-within-GS iteratively alternates the following simulation steps

$$\rho^{(l+1)}, \underline{p}' \sim \text{TJM},$$
$$y_{st}^{(l+1)} | \pi_s^{-1}, \rho^{(l+1)}, \underline{p}' \sim \text{Exp}\left(\sum_{i=1}^{K} \delta_{sti}^{(l+1)} p_i'\right),$$
$$p_i^{(l+1)} | \underline{\pi}^{-1}, \underline{y}^{(l+1)}, \rho^{(l+1)} \sim \text{Ga}\left(c+N, d+\sum_{s=1}^{N}\sum_{t=1}^{K} \delta_{sti}^{(l+1)} y_{st}^{(l+1)}\right).$$

3 EPL diagnostic

Simulation studies confirmed the efficacy of the TJM-within-GS to recover the actual generating EPL, together with the benefits of the SM strategy to speed up the MCMC algorithm in the exploration of the posterior distribution. However, we were surprised to verify a less satisfactory performance of the TJM-within-GS in terms of posterior exploration in the application to some real-world examples, such as the famous `song` dataset analyzed by [2]. Since the joint proposal distribution relies on summary statistics, the posterior sampling procedure is expected to work well as long as the data are actually taken from an EPL distribution. So, the unexpectedly bad behavior of the MCMC suggested to conjecture that, for such real data, the EPL does not represent the true (or in any case an appropriate) data generating mechanism. This has motivated us to the develop some new tools to appropriately check the model mis-specification issue.

Suppose we have some data simulated from an EPL model. We expect the marginal frequencies of the items at the first stage to be ranked according to the order of the corresponding support parameter component. On the other hand, although computationally demanding to be evaluated in terms of their closed form formula we expect the marginal frequencies of the items at the last stage to be ranked according to the reverse order of the corresponding support parameter component. After proving such a statement one can then derive that the ranking of the marginal frequencies of the items corresponding to the first and last stage should sum up to

$(K+1)$, no matter what their support is. Of course, this is less likely to happen when the sample size is small or when the support parameters are not so different of each other. In any case, one can define a test statistic by considering, for each couple of integers (j, j') candidate to represent the first and the last stage ranks, namely $\rho(1)$ and $\rho(K)$, a discrepancy measure $T_{jj'}(\pi)$ between $K+1$ and the sum of the rankings of the frequencies corresponding to the same item extracted in the first and in the last stage. Formally, let $\underline{r}_j^{[1]} = (r_{j1}^{[1]}, \ldots, r_{jK}^{[1]})$ and $\underline{r}_{j'}^{[K]} = (r_{j'1}^{[K]}, \ldots, r_{j'K}^{[K]})$ be the marginal item frequency distributions for the j-th and j'-th positions, to be assigned respectively at the first [1] and last [K] stage. In other words, the generic entry $r_{ji}^{[s]}$ is the number of times that item i is ranked j-th at the s-th stage. The proposed EPL diagnostic relies on the following discrepancy

$$T_{jj'}(\pi) = \sum_{i=1}^{K} |(\text{rank}(\underline{r}_j^{[1]})_i + \text{rank}(\underline{r}_{j'}^{[K]})_i - (K+1))|,$$

implying that the smaller the test statistics, the larger the plausibility that the two integers (j, j') represent the first and the last components of the reference order. To globally assess the conformity of the sample with the EPL, we consider the minimum value of $T_{jj'}(\pi)$ over all the possible rank pairs satisfying the order constraints

$$T(\pi) = \min_{(j,j') \in \mathscr{P}} T_{jj'}(\pi), \tag{2}$$

where $\mathscr{P} = \{(j, j') : j \in \{1, K\} \text{ and } j \neq j'\}$.

3.1 Applications to real data

We fit the EPL with reference order constraints to the `sport` dataset of the `Rankcluster` package, where N=130 students at the University of Illinois were asked to rank K=7 sports in order of preference: 1=Baseball, 2=Football, 3=Basketball, 4=Tennis, 5=Cycling, 6=Swimming and 7=Jogging. We estimated the Bayesian EPL with hyperparameter setting $c = d = 1$, by running the TJM-within-GS for 20000 iterations and discarding the first 2000 samplings as burn-in phase. We show the approximation of the posterior distribution on the reference order in Figure 1, where it is apparent that the MCMC is mixing sufficiently fast and there is some uncertainty on the underlying reference order. The modal reference order is (7,1,2,3,4,6,5), with slightly more than 0.4 posterior probability. However, when we compared the plausibility of the observed diagnostic statistic with the reference distribution under the fitted EPL, we got a warning with a bootstrap classical p-value approximately equal to 0.011. This should indeed cast some doubt on the use of PL or EPL as a suitable model for the entire dataset. In fact, we have verified that, after suitably splitting the dataset into two groups according to the EPL mixture methodology suggested by [4] (best fitting 2-component EPL mixture

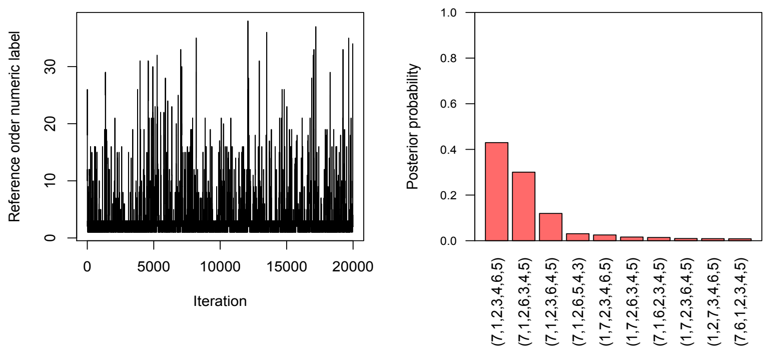

Fig. 1: Traceplot (left) and top-10 posterior probabilities (right) for the reference order parameter.

with BIC=2131.20), we have a different more comfortable perspective for using the EPL distribution to separately model the two clusters. The modal reference orders are (1,2,3,4,5,6,7) and (1,2,3,7,4,5,6) and the estimated Borda orderings are (7,6,4,5,3,1,2) and (1,2,3,4,6,7,5), indicating opposite preferences in the two subsamples towards team and individual sports. In this case, no warning by the diagnostic tests applied separately to the two subsamples is obtained, since the resulting p-values are 0.991 and 0.677.

4 Conclusions

We have addressed some relevant issues in modelling choice behavior and preferences. In particular, we have further explored the idea in [4] related to the use of the reference order specifying the order of the ranks sequentially assigned by introducing monotonicity restrictions on the discrete parameter to describe a "top-or-bottom" attribution of the positions. Our contribution allows to gain more insights on the sequential mechanism of formation of preferences, whether or not it is appropriate at all and whether it privileges a more or less natural ordered assignment of the most extreme ranks. Additionally, some issues experienced when implementing a well-mixing MCMC approximation motivated us to derive a diagnostic tool to test the appropriateness of the EPL distribution, whose effectiveness has been checked with an application to a real example.

References

1. Alvo M, Yu PL (2014). *Statistical methods for ranking data*. Springer.
2. Critchlow DE, Fligner MA, Verducci JS (1991). "Probability models on rankings." *Journal of Mathematical Psychology*, **35**(3), 294–318.
3. Marden JI (1995). *Analyzing and modeling rank data*, volume 64 of *Monographs on Statistics and Applied Probability*. Chapman & Hall. ISBN 0-412-99521-2.
4. Mollica C, Tardella L (2014). "Epitope profiling via mixture modeling of ranked data." *Statistics in Medicine*, **33**(21), 3738–3758. ISSN 0277-6715. doi:10.1002/sim.6224.
5. Mollica C, Tardella L (2017). "Bayesian mixture of Plackett-Luce models for partially ranked data." *Psychometrika*, **82**(2), 442–458. ISSN 0033-3123. doi:10.1007/s11336-016-9530-0.
6. Mollica C, Tardella L (2018). "Algorithms and diagnostics for the analysis of preference rankings with the Extended Plackett-Luce model." *arXiv preprint: http://arxiv.org/abs/1803.02881*.