Probabilistic preference learning with the Mallows rank model

Elja Arjas

Abstract

Ranking and comparing items is crucial for collecting information about preferences in many areas, including marketing and politics. The Mallows rank model is among the most successful approaches to analyse rank data, but its computational complexity has limited its use to a particular form based on the Kendall distance. In the talk, a computationally tractable method for Bayesian inference in Mallows models with any right-invariant metric is presented. The method performs inference on the consensus ranking of the items, also when based on only partial rankings, such as top-k items or pairwise comparisons. Possible heterogeneity of the assessors is handled by introducing a mixture model for clustering the assessors into more homogeneous subgroups, with cluster-specific consensus rankings. Probabilistic predictions on the class membership of assessors based on their ranking of only some items are provided, thereby also enabling the prediction of missing individual preferences. The approach is tested by using several experimental and benchmark datasets. – The talk is based on joint work with Valeria Vitelli, Øystein Sørensen and Arnoldo Frigessi.