Non parametric Bayesian priors for hidden Markov random fields

Florence Forbes

INRIA Grenoble Rhone-Alpes

Abstract. One of the central issues in statistics and machine learning is how to select an adequate model that can automatically adapt its complexity to the observed data. Bayesian nonparametric methods are thought of as one of the most promising candidates that are capable of handling such tasks. In the present work, we consider the issue of determining from the data the number of groups in a clustering task of non independent observations. The required guess on the number of clusters is avoided by considering models with an infinite number of clusters as suggested in Dirichlet Process Mixture models (DPMM). However, for tasks such as unsupervised image segmentation with spatial relationships or dependencies between the observations, DPMM are not satisfying. We propose to combine a Markov random field model with different Bayesian nonparametric priors and illustrate such a combination on a Potts model combined with Dirichlet and Pitman-Yor processes. As regards inference, the variational expectation-maximization algorithm is adopted due to its lower computational cost with respect to its MCMC counterpart. Finally, the proposed framework is illustrated on image segmentation and risk mapping from count data.