
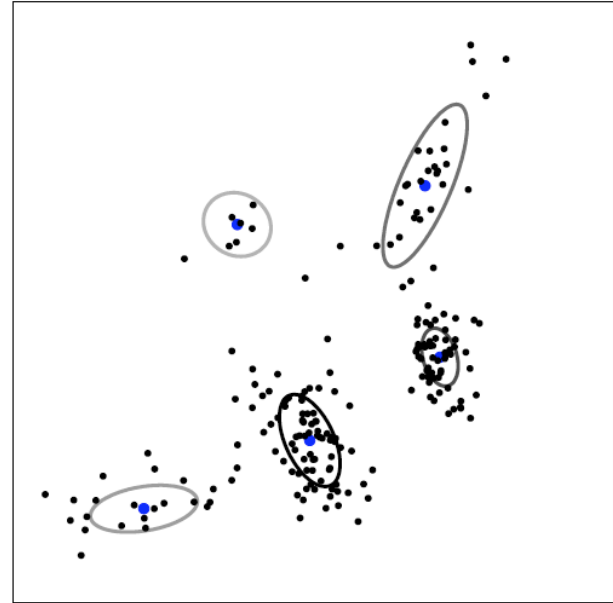




Learning and Inference in Probabilistic Graphical Models

Gibbs Sampling for Exponential Family Mixtures
April 19, 2010


Finite Exponential Family Mixtures



Standard Gibbs: 2 Iterations




$\log p(x | \pi, \theta) = -539.17$




$\log p(x | \pi, \theta) = -497.77$

Standard Gibbs: 10 Iterations




$\log p(x \mid \pi, \theta) = -404.18$




$\log p(x \mid \pi, \theta) = -454.15$

Standard Gibbs: 50 Iterations

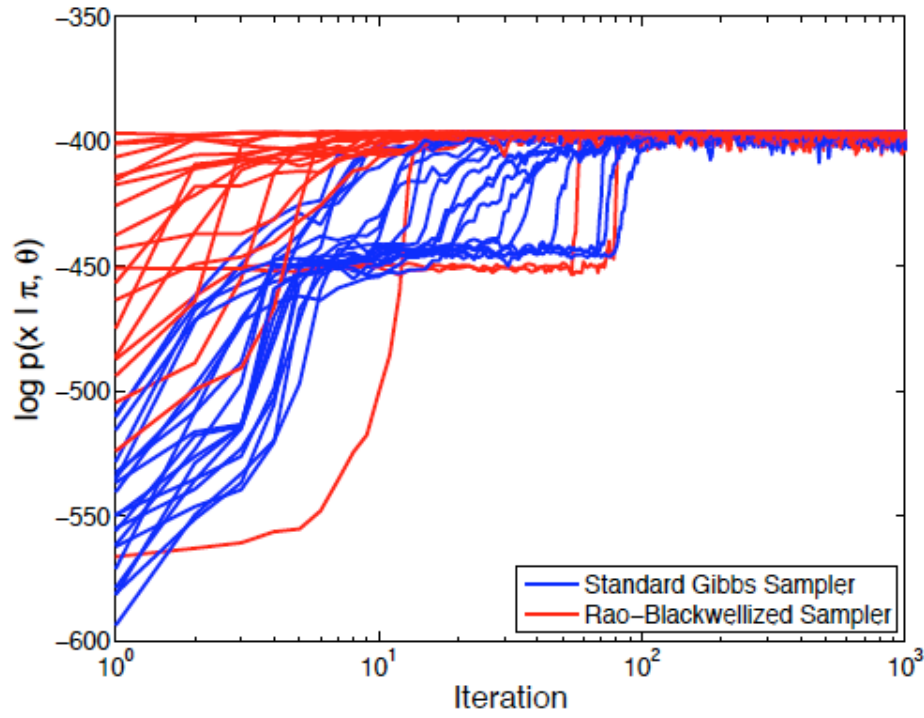


$$\log p(x \mid \pi, \theta) = -397.40$$

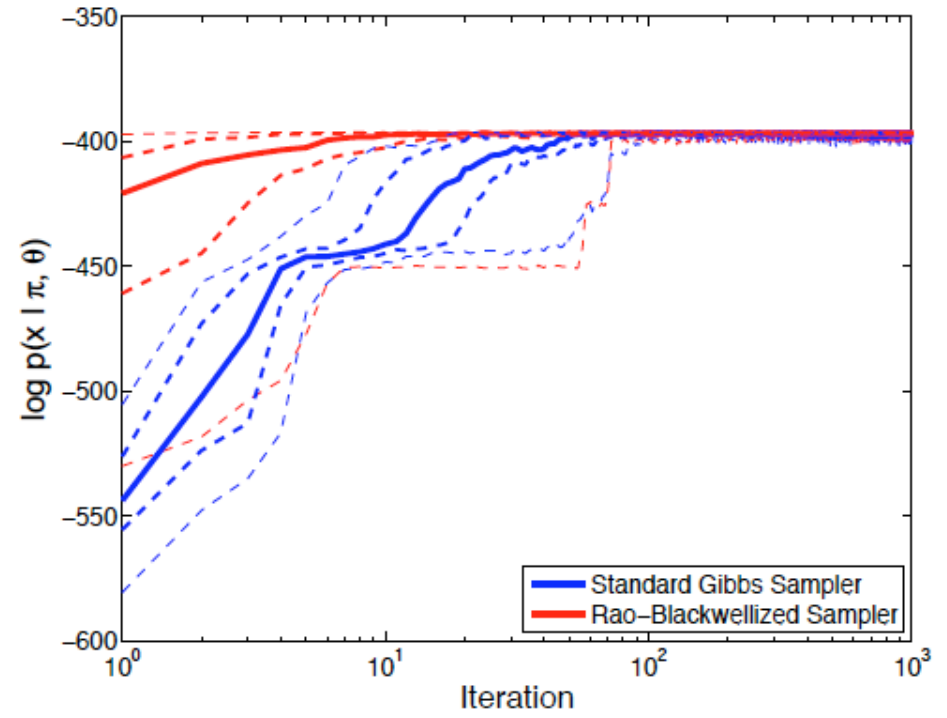


$$\log p(x \mid \pi, \theta) = -442.89$$

Convergence: Multiple Chains

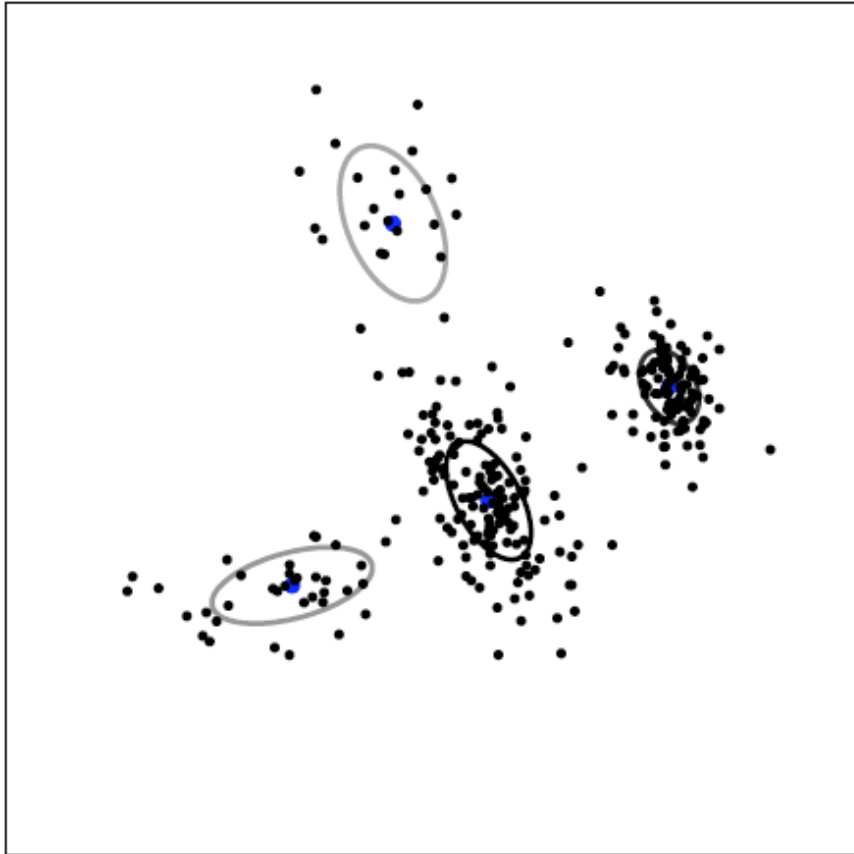


Trajectories for 20 initializations




*Quantiles for 100 initializations:
0.05, 0.25, 0.50, 0.75, 0.95*

Rao-Blackwellized Gibbs: 2 Iterations




$\log p(x \mid \pi, \theta) = -399.06$




$\log p(x \mid \pi, \theta) = -461.94$

Rao-Blackwellized Gibbs: 10 Iterations




$\log p(x \mid \pi, \theta) = -397.38$




$\log p(x \mid \pi, \theta) = -449.23$

Rao-Blackwellized Gibbs: 50 Iterations



$$\log p(x \mid \pi, \theta) = -396.53$$



$$\log p(x \mid \pi, \theta) = -448.68$$