

Image Segmentation with Spatial Priors

Satvik Saha

Department of Statistics, Columbia University

Segmentation with Mixture Models

Let $x_\ell \in \mathbb{R}^d$ be the color of a pixel at location $\ell \in \mathcal{L}$. A K -component mixture model generates an image $\mathbf{x}$ via

$$\begin{aligned} \boldsymbol{\theta} &\sim p(\boldsymbol{\theta}), \\ \boldsymbol{\beta} &\sim p(\boldsymbol{\beta}), \\ \mathbf{z} \mid \boldsymbol{\beta} &\sim p(\mathbf{z} \mid \boldsymbol{\beta}), \\ x_\ell \mid \boldsymbol{\theta}, z_\ell = k &\sim p(x_\ell \mid \theta_k, z_\ell = k), \quad \ell \in \mathcal{L}. \end{aligned}$$

Segmentation involves inferring the distribution of labels $\mathbf{z} \mid \mathbf{x}$.

Naïve Mixture Models

Simply choosing $z_\ell \mid \boldsymbol{\beta} \stackrel{iid}{\sim} p(\cdot \mid \boldsymbol{\beta})$ ignores spatial information, and uses only the observed distribution of color $\mathbf{x}$; shuffling the locations of all pixels would produce the same labels! This can produce rough, grainy segments with ‘islands’.




Fig. 1: A naïve mixture model only sees the distribution of x_ℓ .

We will focus on one particular aspect of spatial information, and use the following heuristic: *neighboring pixels often belong to the same segment*.

Potts-Markov Spatial Priors

Given some description $\mathcal{G}(\mathcal{L}, \mathcal{E})$ of pixel neighborhoods, consider

$$p(\mathbf{z} \mid \boldsymbol{\beta}) = \frac{1}{Z(\boldsymbol{\beta})} \exp \left(\sum_{\ell \sim \ell'} \sum_{k=1}^K \beta_k \mathbf{1}(z_\ell = z_{\ell'} = k) \right) = \frac{\exp(\boldsymbol{\beta}^\top \mathbf{T}(\mathbf{z}))}{Z(\boldsymbol{\beta})}.$$

This rewards assigning neighboring pixels the same label, producing ‘smoother’ segments.




Fig. 2: Draws from $p(\mathbf{z} \mid \boldsymbol{\beta})$, using Gibbs sampling.

Gaussian-GIW-Potts Mixture

Using the Potts-Markov spatial prior, we further choose

$$\begin{aligned} \beta_k &\stackrel{iid}{\sim} \text{Gamma}(a_\beta, b_\beta), \\ \Sigma_k &\stackrel{iid}{\sim} \text{Inv-Wishart}(\Sigma_k; \Sigma_0, \nu_0), \\ \mu_k \mid \Sigma_k &\sim \text{Normal}(m_0, \kappa_0^{-1} \Sigma_k), \\ x_\ell \mid \mu_k, \Sigma_k, z_\ell = k &\sim \text{Normal}(\mu_k, \Sigma_k). \end{aligned}$$



Original




K-Means




GMM



GGIWP-MM, $\beta = 1$



GGIWP-MM



Posterior summary of β

Fig. 3: Image segmentation, with $K = 5$.

The posterior distribution of $\boldsymbol{\beta}$ gives information about the smoothness of each segment.

Auxiliary Variable MCMC

In this model, the joint likelihood itself is intractable. Specifically, it is difficult to draw from the full-conditional

$$p(\boldsymbol{\beta} \mid \mathbf{z}) \propto \frac{p(\boldsymbol{\beta})}{Z(\boldsymbol{\beta})} \exp \left(\boldsymbol{\beta}^\top \mathbf{T}(\mathbf{z}) \right).$$

For instance, a Metropolis-Hastings acceptance ratio depends on the intractable term $Z(\boldsymbol{\beta}^*)/Z(\boldsymbol{\beta})$.

To circumvent this, we use an auxiliary variable exchange algorithm (Møller et al., 2004):

1. Propose $\boldsymbol{\beta}^* \sim q(\boldsymbol{\beta}^*; \boldsymbol{\beta})$.
2. Draw auxiliary $\mathbf{z}^* \sim p(\mathbf{z}^* \mid \boldsymbol{\beta}^*)$.
3. Accept $\boldsymbol{\beta}^*$ with probability

$$\frac{q(\boldsymbol{\beta}; \boldsymbol{\beta}^*) p(\boldsymbol{\beta}^*)}{q(\boldsymbol{\beta}^*; \boldsymbol{\beta}) p(\boldsymbol{\beta})} \cdot \frac{p(\mathbf{z} \mid \boldsymbol{\beta}^*)}{p(\mathbf{z} \mid \boldsymbol{\beta})} \cdot \frac{p(\mathbf{z}^* \mid \boldsymbol{\beta})}{p(\mathbf{z}^* \mid \boldsymbol{\beta}^*)}.$$

This is fairly expensive (Gibbs within Metropolis-Hastings within Gibbs), but yields the correct posterior as its stationary distribution!