Lecture 9: Estimation and segmentation using MRFs

Spatial Statistics and Image Analysis

David Bolin
University of Gothenburg

Gothenburg
May 6, 2019

Example of discrete MRFs

A general MRF for lattices with first-order neighborhoods:

$$p(\mathbf{z}) = \frac{1}{Z}\exp\left(\sum_i \alpha_{z_i} + \frac{1}{2}\sum_i \sum_{j\in N_i} \beta_{z_i,z_j}\right).$$

Here $\{\alpha_1,\ldots,\alpha_K\}$ determines the prior probabilities for the K classes and $\beta_{k,l}$ determines the interaction between class k and class l.

Common simplifications include assuming that

$$\beta_{k,l} = \begin{cases}\beta_k & \text{if } l = k\\ 0 & \text{otherwise}\end{cases} \quad \text{or} \quad \beta_{k,l} = \begin{cases}\beta & \text{if } l = k\\ 0 & \text{otherwise}\end{cases}$$

Markov random field mixture models

- Hierarchical model for pixel values given classes:

$$\pi(\mathbf{Y}_i|z_i = k) \sim \mathsf{N}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$$

$$\pi(z_i) = \begin{cases}\pi_1 & \text{if } z_i = 1\\ \pi_2 & \text{if } z_i = 2\\ \vdots & \\ \pi_K & \text{if } z_i = K\end{cases}$$

- Assuming independence between the pixels is not realistic!
- In a Markov random field mixture model, we use the model

$$\pi(\mathbf{Y}_i|z_i = k) \sim \mathsf{N}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$$
$$\mathbf{z} \sim \pi(\mathbf{z})$$

  here $\mathbf{z} = (z_1,\ldots,z_n)$ is a random field that takes values in $\{1,\ldots,K\}$, with density $\pi(\mathbf{z})$.
- Spatial dependencies modeled through $\pi(\mathbf{z})$.

Conditional distributions and sampling

- The normalizing constant Z is intractable. However, the conditional distributions are simple:

$$p(z_i|\mathbf{z}_{-i}) = \frac{\exp(\alpha_{z_i} + \beta\sum_{j\in N_i} 1(z_j = z_j))}{\sum_k \exp(\alpha_k + \beta\sum_{j\in N_i} 1(z_j = k))}.$$

- Since we have simple conditional distributions, we can sample the field using Gibbs sampling:

1. Choose a starting value $\mathbf{z}^{(0)}$.
2. Repeat for $j = 1,\ldots,N$:
   - Sample $z_1^{(j)}$ from $\pi(z_1|z_2^{(j-1)},\ldots,z_n^{(j-1)})$.
   - Sample $z_2^{(j)}$ from $\pi(z_2|z_1^{(j)},z_3^{(j-1)},\ldots,z_n^{(j-1)})$.

     $\vdots$

   - Sample $z_n^{(j)}$ from $\pi(z_n|z_1^{(j)},\ldots,z_{n-1}^{(j)})$.
3. $\mathbf{z}^{(K)},\ldots,\mathbf{z}^{(N)}$ are dependent draws from $\approx \pi(\mathbf{z})$.

Parallell Gibbs sampling

- For an MRF with a first-order neighborhood structure, the black nodes are conditionally independent given the white.
- This means that we can do the Gibbs sampling in parallell:

1. Choose a starting value $\mathbf{z}^{(0)}$.
2. Repeat for $j = 1, \ldots, N$:
   - Sample $\mathbf{z}_{white}^{(j)}$ from $\pi(\mathbf{z}_{white}|\mathbf{z}_{black}^{(j-1)})$.
   - Sample $\mathbf{z}_{black}^{(j)}$ from $\pi(\mathbf{z}_{black}|\mathbf{z}_{white}^{(j)})$.
3. $\mathbf{z}^{(K)}, \ldots, \mathbf{z}^{(N)}$ are dependent draws from $\approx \pi(\mathbf{z})$.

Sampling each element in $\mathbf{z}_{white}^{(j)}$ is done as in the previous sampler.

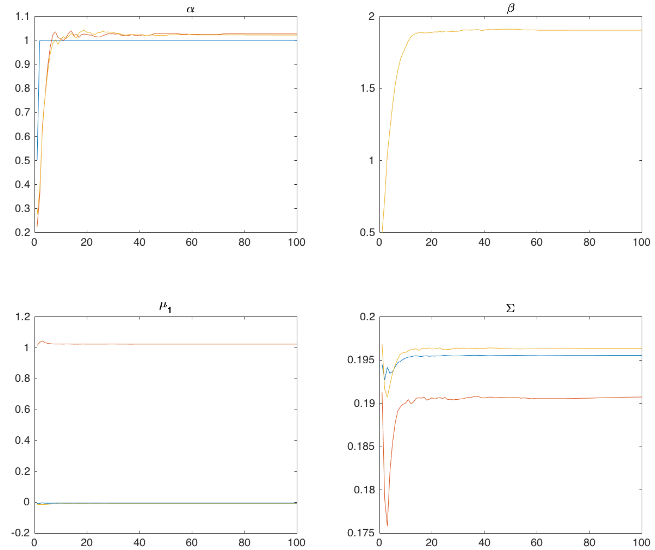

Example data

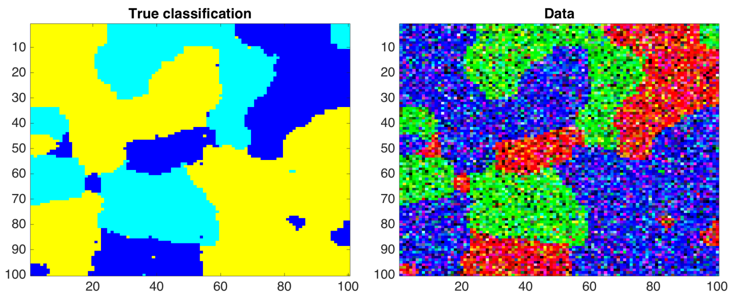

Parameters:

$$\alpha_k = \frac{1}{3} \qquad \beta = 2 \qquad \mathbf{\Sigma}_k = 0.2\mathbf{I}$$

$$\boldsymbol{\mu}_1 = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \qquad \boldsymbol{\mu}_2 = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} \qquad \boldsymbol{\mu}_3 = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$$