

Exercise: Gibbs sampling for a 1D Gaussian mixture model

Consider applying Gibbs sampling to a univariate mixture of Gaussians. Derive the expressions for the full conditionals. Hint: if we know $z_n = j$ (say), then μ_j gets “connected” to x_n , but all other values of μ_i , for all $i \neq j$, are irrelevant. (This is an example of context-specific independence, where the structure of the graph simplifies once we have assigned values to some of the nodes.) Hence, given all the z_n values, the posteriors of the μ ’s should be independent, so the conditional of μ_j should be independent of μ_{-j} . (Similarly for σ_j .)