The file ('Data for exercise 6') contains data sampled from a mixture of $K = 4$ Gaussians. The standard deviations of the Gaussians are known and identical, $\sigma_k = \sigma = 0.5$. Reconstruct the unknown parameters and latent variables,

$$p_k, \mu_k, Z_i$$

where Z_i represents the assignment of each point to one of the mixture components. To perform the inference, use Gibbs sampling:

- Initialize randomly p_k, μ_k, Z_i. Compute $N_k = \sum_i \chi(Z_i = k), m_k = \sum_i \chi(Z_i = k) x_k$.
- Perform Gibbs sampling iterations
  - Sample μ_k from a normal with mean $\mu'_k = \left(\frac{\mu_0}{\sigma_0^2} + \frac{m_k}{\sigma^2}\right) \frac{1}{\left(\frac{1}{\sigma_0^2} + \frac{N_k}{\sigma^2}\right)}$ and std. dev. $\sigma'_k = \left(\frac{1}{\sigma_0^2} + \frac{N_k}{\sigma^2}\right)^{-1/2}$. Here $\mu_0 = 0$ and $\sigma_0 = 1000$ are the Gaussian prior parameters. Do this for all k.
  - Sample the p from a Dirichlet distribution with parameters $\gamma'_k = \gamma_k + N_k$. Here, $\gamma_k = 1$ are the Dirichlet prior parameters.
  - Sample the Z_i from a categorical disribution, $Prob(Z_i = k) = \frac{p_k e^{-\frac{(x_k - \mu_k)^2}{2\sigma^2}}}{\sum_k p_k e^{-\frac{(x_k - \mu_k)^2}{2\sigma^2}}}$. If Z_i is updated, update N_k, m_k.