Gibbs Sampling

A Markov Chain algorithm to sample from the joint distribution $f_{\mathbf{X}}$ of a p-dimensional vector of random variables $\mathbf{X} = (X_1, \ldots, X_p)$.

1. Specify an initial value $\mathbf{X}^{(0)} = \left(X_1^{(0)}, \ldots, X_p^{(0)}\right)$
2. Repeat for $j = 1, 2, \ldots,$.
   (a) Generate $X_1^{(j+1)}$ from $f\left(X_1 | X_2^{(j)}, X_3^{(j)}, \ldots, X_p^{(j)}\right)$
   (b) Generate $X_2^{(j+1)}$ from $f\left(X_2 | X_1^{(j+1)}, X_3^{(j)}, \ldots, X_p^{(j)}\right)$
   (c) $\vdots$
   (d) Generate $X_p^{(j+1)}$ from $f\left(X_p | X_1^{(j+1)}, \ldots, X_{p-1}^{(j+1)}\right)$.
3. Return the values $\mathbf{X^{(1)}}, \mathbf{X^{(2)}}, \mathbf{X^{(3)}}, \ldots,$.

This algorithm is a markov chain. The values that the algorithm generates are realisations of the random vector $\mathbf{X}$. The state space of the MC is therefore the set of values that $\mathbf{X}$ can take.

After running the MC for a while, it will converge and the values of $\mathbf{X}$ generated will be a random sample from the limiting distribution $f_{\mathbf{X}}$.