Gibbs Sampler

- Way of simulating from multivariate distributions based on the ability to simulate from conditional distn's
  - → i.e. appropriate when sampling from marginal distns is not convenient or possible

- Sampling from bivariate distributions:

  i.e. consider a bivariate normal distribution

$$f(x,y) = f(x)\, f(y|x)$$

  So we can simulate from the joint density $f(x,y)$ by first simulating $X=x$ from $f(x)$ & then simulating $Y=y$ from $f(y|x)$.

- Gibbs sampler is similar...

Sequence of samples can be used to approximate the joint distribution, the marginal dist'n of one of the variables, or some subset of the variables, or to compute approx. expected value

↑ e.g. unknown parameters

- Gibbs sampling is commonly used for statistical inference (i.e. Bayesian inference)

- It's an alternative to deterministic algorithms for statistical inference such as the expectation-maximization algorithm (EM).

  ↑ randomized algorithm / makes use of random #'s

- Applicable when the joint dist'n is not known explicitly (or difficult to sample from directly), but the conditional dist'n of each variable is known & easier to sample from.

  → Gibbs sampling algorithm generates a sample from the dist'n of each variable conditional on the current values of the other variables. Seq. of samples constitutes a MC w/ desired joint dist'n.

  e.g. sampling the posterior dist'n of a Bayesian network

See R code: mcmc-examples-solns.r
gibbs-sampler-examples.r