STAT 535

MCMC Continued

- Acceptance Rates and Goldilocks
- How troublesome is dependent sampling
- Univariate versus multivariate approaches

Acceptance Rates and Goldilocks

Recall in the toy example we used *random walk* MH:

$$Z^*|Z^{(i)} \quad \sim \quad N(Z^{(i)}, \tau^2)$$

Choice of τ?

Intuitively, acceptance rate *decreases* as τ increases (why)?

Would an acceptance rate close to 0% or close to 100% be good?

Goldilocks: a 'mid-range' acceptance rate is just right!

Toy problem revisited:

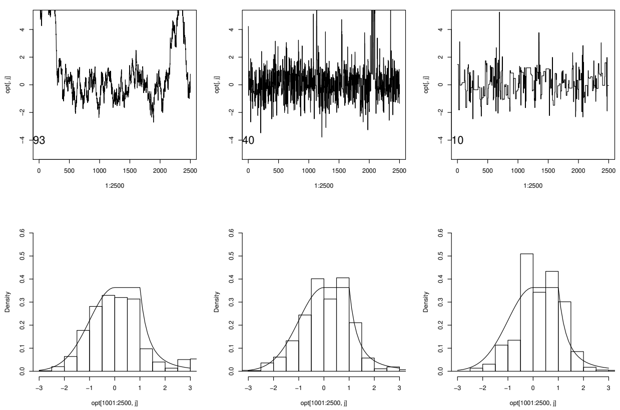

Dependence in MCMC Output - A Problem?

Say $Z^{(1)}, \ldots, Z^{(m)}$ are 'SDID' as π. For 'learning' about π (e.g. computing $E_\pi\{g(Z)\}$), is the dependence a problem?

NO, in that (under conditions) a LLN still applies. As $m \to \infty$,

$$m^{-1} \sum_{i=1}^{m} g(Z^{(i)}) \quad \to \quad E_\pi\{g(Z)\}.$$

YES, in that

$$Var\left\{m^{-1} \sum_{i=1}^{m} g(Z^{(i)})\right\} \quad \approx \quad \frac{\sigma^2}{m}\left(1 + 2\sum_{k=1}^{\infty} \rho_k\right)$$

where $\sigma^2 = Var_\pi\{g(Z)\}$ and $\rho_k = Corr(g(Z^{(i)}), g(Z^{(i+k)}))$ (which doesn't depend on i).

Multivariate versus Univariate MCMC

Say $Z = (Z_1, Z_2, Z_3)$.

Multivariate update: Obtain $z^{(i+1)}$ from $z^{(i)}$ by MH with $\pi(Z)$ as the target distribution (whole vector either accepted or rejected).

Univariate update:

1. Obtain $z_1^{(i+1)}$ from $z_1^{(i)}$ by MH,
$\pi\left(Z_1 | Z_2 = z_2^{(i)}, Z_3 = z_3^{(i)}\right)$ is target distribution.

2. Obtain $z_2^{(i+1)}$ from $z_2^{(i)}$ by MH,
$\pi\left(Z_2 | Z_1 = z_1^{(i+1)}, Z_3 = z_3^{(i)}\right)$ is target distribution.

3. Obtain $z_3^{(i+1)}$ from $z_3^{(i)}$ by MH,
$\pi\left(Z_3 | Z_1 = z_1^{(i+1)}, Z_2 = z_2^{(i+1)}\right)$ is target distribution.

Reduce a high-D problem to many one-D problems!

In fact, early interest in MCMC for Bayesian applications arose from problems where $\pi(z_1, \ldots, z_p)$ was intractable, but each *full conditional* was a standard distribution that could be sampled from easily. That is,

$$\pi(z_i | z_{(i)}) \quad = \quad \pi(z_i | z_1, \ldots, z_{i-1}, z_{i+1}, \ldots, z_p)$$

is a standard distribution for each i.

Then one can simply sample from each full conditional in turn. This is the *Gibbs Sampler* — special case of MH with univariate updates, acc. prob. = 1.

Proportionality idea: Often can stare at the unnormalized joint density and identify what the full conditionals must be.