Markov chains

1. Let $(X_n)_{n\geq 0}$ be a stationary Markov chain. Show that $\mathrm{Cov}(X_k, X_\ell)$ depends only on $|k-\ell|$.

2. Let $(\varepsilon_n)_{n\geq 0}$ be i.i.d. Gaussian variables with zero mean and variance σ^2. Let $|\alpha| < 1$ and consider the AR(1) process

$$X_n = \alpha X_{n-1} + \varepsilon_n, \quad n \geq 1,$$

with $X_0 = \varepsilon_0$.

   (a) Find the mean and the variance of X_n. Is $(X_n)_{n\geq 0}$ stationary?

   (b) Show that for all $0 \leq h \leq n$,

$$\mathrm{Corr}(X_n, X_{n-h}) = \alpha^h \sqrt{\frac{\mathrm{Var}(X_{n-h})}{\mathrm{Var}(X_n)}}.$$

   (c) Show that $\lim_{n\to\infty} \mathrm{Var}(X_n) = \sigma^2/(1-\alpha^2)$ and $\lim_{n\to\infty} \mathrm{Corr}(X_n, X_{n-h}) = \alpha^h$.

   (d) Now, suppose that $X_0 = \varepsilon_0/\sqrt{1-\alpha^2}$. Is $(X_n)_{n\geq 0}$ stationary?

The Gibbs sampler

3. Let p and q be Markov transition densities on $\mathsf{X} \subseteq \mathbb{R}^d$. The *product* $p \otimes q$ of p and q is defined as

$$[p \otimes q](z \mid x) = \int p(y \mid x) q(z \mid y)\, dy \quad ((x,z) \in \mathsf{X}^2).$$

   (a) Show that $p \otimes q$ is a Markov transition kernel on X.

   (b) Assume that p and q both allow π as a stationary distribution. Show that also $p \otimes q$ allows π as a stationary distribution.

4. Recall that the Gibbs sampler simulates an m-variate Markov chain $(X_n)_{n\geq 0}$ having some multivariate distribution π as stationary distribution by, in each sub-step, sampling from the conditional distributions $\pi_\ell(x^\ell \mid x^{-\ell})$, where $x^{-\ell} = (x^1, \ldots, x^{\ell-1}, x^{\ell+1}, \ldots, x^m)$.

   (a) Show that each sub-step of the Gibbs sampler is π-reversible (i.e., satisfies detailed balance for π).

   (b) Show that one full iteration (comprising m sub-steps) of the Gibbs sampler (see Lecture 10) allows π as a stationary distribution.

Barker's MCMC algorithm

5. *Barker's MCMC algorithm* targeting some density π (known up to a normalizing constant) on $\mathsf{X} \subseteq \mathbb{R}^d$ generates a Markov chain $(X_n)_{n\geq 0}$ as follows: given X_n,

   draw $X^* \sim r(x \mid X_n)$;

   set $X_{n+1} \leftarrow \begin{cases} X^* & \text{w. pr. } \frac{\pi(X^*)}{\pi(X^*)+\pi(X_n)} \\ X_n & \text{otherwise} \end{cases}$;

   Here r is some *symmetric* proposal transition density.

   (a) Find the transition density q of Barker's algorithm.

   (b) Show that q is π-reversible.

(c) What should be regarded as a main difference between Barker's algorithm and the Metropolis-Hastings algorithm (with symmetric proposal distribution r)? Which method seems preferable and why? (Hint: consider the situation where $\pi(X^*)$ is very close to $\pi(X_n)$.)