1 Choose-your-own-adventure (1-d Ising model).

The Ising model is a spatial process on the 1-d lattice, inspired by ferromagnets in chemistry. There are N spins denoted $\boldsymbol{X} = (X_1, \ldots, X_N)$, and each spin is assigned either $X_i = -1$ or $X_i = +1$. Assuming "free" boundary conditions, the probability mass function for vectors $\boldsymbol{X}$ is

$$p(\boldsymbol{x}) = \frac{1}{Z_\beta} \exp\bigl(\beta \sum\nolimits_{i=1}^{N-1} x_i x_{i+1}\bigr), \qquad Z_\beta = \sum_{\boldsymbol{x} \in \{-1,+1\}^N} \exp\bigl(\beta \sum\nolimits_{i=1}^{N-1} x_i x_{i+1}\bigr),$$

where $\beta > 0$ is the inverse temperature parameter.

(a) There is an exact sampler for the 1-d Ising model. First, define functions $V_1(x_1) = 1$ and $V_{t+1}(x_{t+1}) = \sum_{x_t \in \{-1,+1\}} V_t(x_t) \exp(\beta x_t x_{t+1})$ for $t = 1, \ldots, N-1$ and calculate the partition function $Z_\beta = \sum_{x_N \in \{-1,+1\}} V_N(x_N)$. Next, sample the last spin X_N from the distribution $p_N(x_N) = V_N(x_N)/Z$, and recursively sample each spin X_{t-1} from the distribution

$$p_{t-1}(x_{t-1}) = \frac{V_{t-1}(x_{t-1}) \exp(\beta x_{t-1} X_t)}{\sum_{y_{t-1} \in \{-1,+1\}} V_{t-1}(y_{t-1}) \exp(\beta y_{t-1} X_t)}$$

for $t = N, N-1, \ldots, 2$. Try this on a computer, or write down a simple analytic formula.

(b) Use part (a) to calculate the correlation function $C(t) = \mathbb{E}\bigl[X_0 X_t\bigr]$. How does the correlation function change as we increase β?

2 Choose-your-own-adventure (1-d Gaussian random field).

Consider a Gaussian random field on the 1-d lattice. There are N random variables denoted $\boldsymbol{X} = (X_1, \ldots, X_N)$, and we assume $\boldsymbol{X} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{\Sigma})$ has a multivariate Gaussian distribution.

(a) The first sampling method is to calculate a square root decomposition $\boldsymbol{\Sigma} = \boldsymbol{A}\boldsymbol{A}^*$ (e.g., by Cholesky decomposition) and set $\boldsymbol{X} = \boldsymbol{A}\boldsymbol{Z}$ where $\boldsymbol{Z} = (Z_1, \ldots, Z_N)$ are independent standard Gaussians. Write down or code up how to do this for the moving average process with covariance matrix $\boldsymbol{\Sigma}_{ij} = 1/10 \max\{\min\{i, j, 10 - |i - j|\}, 0\}$.

(b) The second sampling method is to take a square root decomposition of the inverse covariance $\boldsymbol{\Sigma}^{-1} = \boldsymbol{B}\boldsymbol{B}^*$ and solve the linear system $\boldsymbol{B}\boldsymbol{X} = \boldsymbol{Z}$, where $\boldsymbol{Z} = (Z_1, \ldots, Z_N)$ are independent standard Gaussians. Write down or code up how to do this for the autoregressive process with covariance matrix $\boldsymbol{\Sigma}_{i,j} = (9/10)^{|i-j|}$. What differences do you see between the moving average and autoregressive processes?

3 Coding (Poisson point process).

Consider two strategies for sampling from a Poisson point process on the unit square $E = [0, 1]^2$ with intensity function $\lambda(x_1, x_2) = 300(x_1^2 + x_2^2)$. Which strategy is more efficient?

(a) First, there is the "direct" method. Sample the number of points $N \sim \text{Poi}(\int_E \lambda(\boldsymbol{x})\, d\boldsymbol{x})$. Then sample $X_1, \ldots, X_N$ independently from the density $f(\boldsymbol{x}) = \lambda(\boldsymbol{x}) / \int_E \lambda(\boldsymbol{y})\, d\boldsymbol{y}$ (hint: use rejection sampling).

(b) Alternatively, there is the “thinning” method. Generate a homogenous Poisson point process with intensity $\lambda^* = 600$ and thin the points by accepting each point with probability $\lambda(\boldsymbol{x})/\lambda^*$.

4 Coding (Pretty pictures).

In a Matérn process, we first sample centers C from a homogeneous Poisson process on the unit square $E = [0,1]^2$ and then we sample points from a Poisson process on E with intensity function

$$\lambda(\boldsymbol{x}) = \alpha \sum\nolimits_{\boldsymbol{c} \in \mathsf{C}} \mathbb{1}\{\|\boldsymbol{x} - \boldsymbol{c}\| \leq r\}.$$

Play around with parameters until you get pretty pictures. Is the process attractive or repulsive? Can you think of any repulsive point processes?

5 Coding (Poisson hyperplane process)

Let $\Phi = \{t_1, t_2, \ldots\}$ be a Poisson point process on $\mathbb{R}$ with intensity λ and let $\{\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots\}$ be independent random vectors uniformly distributed on the unit circle. Simulate all the hyperplanes $H(\boldsymbol{x}_i, t_i) = \{\boldsymbol{y} \in \mathbb{R}^2 : \langle \boldsymbol{y}, \boldsymbol{x_i} \rangle = t_i\}$ that hit the unit circle. How many hyperplanes are there?

6 Math (Thinning works).

The Laplace functional of a d-dimensional point process Φ is defined for nonnegative functions f by the formula

$$L_\Phi[f] = \mathbb{E}\big[\exp\big(-\sum\nolimits_{\boldsymbol{x} \in \Phi} f(\boldsymbol{x})\big)\big].$$

(a) Calculate the Laplace functional for a Poisson point process with intensity measure Λ.

(b) For any function $p : \mathbb{R}^d \to [0,1]$, a p-thinning of a point process Φ deletes each point $\boldsymbol{x} \in \Phi$ with probability $p(\boldsymbol{x})$. Argue that the p-thinning of a point process is also a point process, and calculate the intensity measure. You can assume: if two point processes share the same Laplace functional, they are the same process (in distribution).

7 Math (Circulants).

Consider a Gaussian random field on a 1-d lattice with N sites and periodic boundary conditions. The covariance matrix takes the form $\boldsymbol{\Sigma}_{i,j} = C(|i-j|)$ for a function C satisfying $C(N-x) = C(x)$, which means that $\boldsymbol{\Sigma}$ is a circulant matrix. We know that a circulant matrix has eigenvectors $1/\sqrt{N}(1, \omega^j, \ldots, \omega^{Nj-j})$ and eigenvalues $\lambda_j = C(0) + C(1)\omega^j + \cdots + C(N-1)\omega^{Nj-j}$, where $\omega = \exp(2\pi i/N)$ is the Nth root of unity. Use this fact to design a slick algorithm for sampling from the $\mathcal{N}(\boldsymbol{0}, \boldsymbol{\Sigma})$ distribution. How could you extend the algorithm to 1-d Gaussian random fields with free boundary conditions, or 2-d Gaussian random fields?