

Bayesian modelling

Markov chain Monte Carlo: Gibbs sampling

Léo Belzile

Last compiled Monday Feb 10, 2025

Reminder: Metropolis–Hastings algorithm

Starting from an initial value $\boldsymbol{\theta}_0$:

1. draw a proposal value $\boldsymbol{\theta}_t^* \sim q(\boldsymbol{\theta} \mid \boldsymbol{\theta}_{t-1})$.
2. Compute the acceptance ratio

$$R = \frac{p(\boldsymbol{\theta}_t^*)}{p(\boldsymbol{\theta}_{t-1})} \frac{q(\boldsymbol{\theta}_{t-1} \mid \boldsymbol{\theta}_t^*)}{q(\boldsymbol{\theta}_t^* \mid \boldsymbol{\theta}_{t-1})}$$

3. With probability $\min\{R, 1\}$, accept the proposal and set $\boldsymbol{\theta}_t \leftarrow \boldsymbol{\theta}_t^*$, otherwise set the value to the previous state, $\boldsymbol{\theta}_t \leftarrow \boldsymbol{\theta}_{t-1}$.

Calculations

We compute the log of the acceptance ratio, $\ln R$, to avoid numerical overflow, with the log posterior difference

$$\ln \left\{ \frac{p(\boldsymbol{\theta}_t^*)}{p(\boldsymbol{\theta}_{t-1})} \right\} = \ell(\boldsymbol{\theta}_t^*) + \ln p(\boldsymbol{\theta}_t^*) - \ell(\boldsymbol{\theta}_{t-1}) - \ln p(\boldsymbol{\theta}_{t-1})$$

Compare the value of $\ln R$ (if less than zero) to $\log(U)$, where $U \sim \text{unif}(0, 1)$.

What proposal?

The *independence* Metropolis–Hastings uses a global proposal q which does not depend on the current state (typically centered at the MAP)

This may be problematic with multimodal targets.

The Gaussian random walk takes $\boldsymbol{\theta}_t^* = \boldsymbol{\theta}_{t-1} + \sigma_p Z$, where $Z \sim \text{Gauss}(0, 1)$ and σ_p is the proposal standard deviation. *Random walks* allow us to explore the space.

Burn in

We are guaranteed to reach stationarity with Metropolis–Hastings, but it may take a large number of iterations...

One should discard initial draws during a **burn in** or warmup period if the chain has not reached stationarity. Ideally, use good starting value to reduce waste.

We can also use the warmup period to adapt the variance of the proposal.

Goldilock principle and proposal variance

Mixing of the chain requires just the right variance (not too small nor too large).




Figure 1: Example of traceplot with proposal variance that is too small (top), adequate (middle) and too large (bottom).

Correlograms for Goldilock




Figure 2: Correlogram for the three Markov chains.

Tuning Markov chain Monte Carlo

- Outside of starting values, the variance of the proposal has a huge impact on the asymptotic variance.
- We can adapt the variance during warmup by increasing/decreasing proposal variance (if acceptance rate is too large/small).
- We can check this via the acceptance rate (how many proposals are accepted).

Optimal acceptance rates

The following rules were derived for Gaussian targets under idealized situations.

- In 1D, rule of thumb is an acceptance rate of 0.44 is optimal, and this ratio decreases to 0.234 when $D \geq 2$ (Sherlock, 2013) for random walk Metropolis–Hastings.
- Proposals for D -variate update should have proposal variance of roughly $(2.38^2 / d) \times \Sigma$, where Σ is the posterior variance.
- For MALA (see later), we get 0.574 rather than 0.234

Block update or one parameter at a time?

As with any accept-reject, proposals become inefficient when the dimension D increase.

This is the **curse of dimensionality**.

Updating parameters in turn

- increases acceptance rate (with clever proposals),
- but also leads to more autocorrelation between parameters

Solutions for strongly correlated coefficients

- Reparametrize the model to decorrelate variables (orthogonal parametrization).
- Block updates: draw correlated parameters together
 - using the chain history to learn the correlation, if necessary

Parameter transformation

Parameters may be bounded, e.g. $\theta_i \in [a, b]$.

- We can ignore this and simply discard proposals outside of the range, by setting the log posterior at $-\infty$ outside $[a, b]$
- We can do a transformation, e.g., $\log \theta_i$ if $\theta_i > 0$ and perform a random walk on the unconstrained space: don't forget Jacobians for $q(\cdot)$!
- Another alternative is to use truncated proposals (useful with more complex algorithms like MALA)

Efficient proposals: MALA

The Metropolis-adjusted Langevin algorithm (MALA) uses a Gaussian random walk proposal

$$\boldsymbol{\theta}_t^* \sim \text{Gauss}\{\mu(\boldsymbol{\theta}_{t-1}), \tau^2 \mathbf{A}\},$$

with mean

$$\mu(\boldsymbol{\theta}_{t-1}) = \boldsymbol{\theta}_{t-1} + \mathbf{A}\eta \nabla \log p(\boldsymbol{\theta}_{t-1} \mid \mathbf{y}),$$

and variance $\tau^2 \mathbf{A}$, for some mass matrix \mathbf{A} , tuning parameter $\tau > 0$.

The parameter $\eta < 1$ is a learning rate. This is akin to a Newton algorithm, so beware if you are far from the mode (where the gradient is typically large)!

Higher order proposals

For a single parameter update θ , a Taylor series expansion of the log posterior around the current value suggests using as proposal density a Gaussian approximation with (Rue & Held, 2005)

- mean $\mu_{t-1} = \theta_{t-1} - f'(\theta_{t-1})/f''(\theta_{t-1})$ and
- precision $\tau^{-2} = -f''(\theta_{t-1})$

We need $f''(\theta_{t-1})$ to be negative!

This gives **local adaption** relative to MALA (global variance).

Higher order and moves

For MALA and cie., we need to compute the density of the proposal also for the reverse move for the expansion starting from the proposal $\mu(\boldsymbol{\theta}_t^*)$.

These methods are more efficient than random walk Metropolis–Hastings, but they require the gradient and the hessian (can be obtained analytically using autodiff, or numerically).

Modelling individual headlines of Upworthy example

The number of conversions `nclick` is binomial with sample size $n_i = \text{nimpression}$.

Since n_i is large, the sample average `nclick/nimpression` is approximately Gaussian, so write

$$Y_i \sim \text{Gauss}(\mu, \sigma^2/n_i)$$

$$\mu \sim \text{trunc. Gauss}(0.01, 0.1^2, 0, 1)$$

$$\sigma \sim \text{expo}(0.7)$$

MALA: data set-up

```
1 data(upworthy_question, package = "hecbayes")
2 # Select data for a single question
3 qdata <- upworthy_question |>
4   dplyr::filter(question == "yes") |>
5   dplyr::mutate(y = clicks/impressions,
6                 no = impressions)
```

MALA: define functions

```
1 # Create functions with the same signature (...) for the algorithm
2 logpost <- function(par, data, ...){
3   mu <- par[1]; sigma <- par[2]
4   no <- data$no
5   y <- data$y
6   if(isTRUE(any(sigma <= 0, mu < 0, mu > 1))){
7     return(-Inf)
8   }
9   dnorm(x = mu, mean = 0.01, sd = 0.1, log = TRUE) +
10  dexp(sigma, rate = 0.7, log = TRUE) +
11  sum(dnorm(x = y, mean = mu, sd = sigma/sqrt(no), log = TRUE))
12 }
```

MALA: compute gradient of log posterior

```
1 logpost_grad <- function(par, data, ...){  
2   no <- data$no  
3   y <- data$y  
4   mu <- par[1]; sigma <- par[2]  
5   c(sum(no*(y-mu))/sigma^2 - (mu - 0.01)/0.01,  
6     -length(y)/sigma + sum(no*(y-mu)^2)/sigma^3 - 0.7  
7   )  
8 }
```

MALA: compute maximum a posteriori

```
1 # Starting values - MAP
2 map <- optim(
3   par = c(mean(qdata$y), 0.5),
4   fn = function(x){-logpost(x, data = qdata)},
5   gr = function(x){-logpost_grad(x, data = qdata)},
6   hessian = TRUE,
7   method = "BFGS")
8 # Check convergence
9 logpost_grad(map$par, data = qdata)
```

MALA: starting values and mass matrix

```
1 # Set initial parameter values
2 curr <- map$par
3 # Compute a mass matrix
4 Amat <- solve(map$hessian)
5 # Cholesky root - for random number generation
6 cholA <- chol(Amat)
```

MALA: containers and setup

```
1 # Create containers for MCMC
2 B <- 1e4L # number of iterations
3 warmup <- 1e3L # adaptation period
4 npar <- 2L
5 prop_sd <- rep(1, npar) # tuning parameter
6 chains <- matrix(nrow = B, ncol = npar)
7 damping <- 0.8
8 acceptance <- attempts <- 0
9 colnames(chains) <- names(curr) <- c("mu", "sigma")
10 # Proposal variance proportional to inverse hessian at MAP
11 prop_var <- diag(prop_sd) %*% Amat %*% diag(prop_sd)
```

MALA: sample proposal with Newton step

```
1 for(i in seq_len(B + warmup)){
2   ind <- pmax(1, i - warmup)
3   # Compute the proposal mean for the Newton step
4   prop_mean <- c(curr + damping *
5     Amat %*% logpost_grad(curr, data = qdata))
6   # prop <- prop_sd * c(rnorm(npar) %*% cholA) + prop_mean
7   prop <- c(mvtnorm::rmvnorm(
8     n = 1,
9     mean = prop_mean,
10    sigma = prop_var))
11 # [...]
```

MALA: reverse step

```
1  # Compute the reverse step
2  curr_mean <- c(prop + damping *
3    Amat %*% logpost_grad(prop, data = qdata))
4  # log of ratio of bivariate Gaussian densities
5  logmh <- mvtnorm::dmvnorm(
6    x = curr, mean = prop_mean,
7    sigma = prop_var,
8    log = TRUE) -
9    mvtnorm::dmvnorm(
10   x = prop,
11   mean = curr_mean,
12   sigma = prop_var,
13   log = TRUE) +
14  logpost(prop, data = qdata) -
15  logpost(curr, data = qdata)
```


MALA: Metropolis–Hastings ratio

```
1  if(logmh > log(runif(1))){  
2    curr <- prop  
3    acceptance <- acceptance + 1L  
4  }  
5  attempts <- attempts + 1L  
6  # Save current value  
7  chains[ind,] <- curr
```

MALA: adaptation

```
1  if(i %% 100 & i < warmup){
2    # Check acceptance rate and increase/decrease variance
3    out <- hecbayes::adaptive(
4      attempts = attempts, # counter for number of attempts
5      acceptance = acceptance,
6      sd.p = prop_sd, #current proposal standard deviation
7      target = 0.574) # target acceptance rate
8    prop_sd <- out$sd # overwrite current std.dev
9    acceptance <- out$acc # if we change std. dev, this is set to zero
10   attempts <- out$att # idem, otherwise unchanged
11   prop_var <- diag(prop_sd) %*% Amat %*% diag(prop_sd)
12 }
13 } # End of MCMC for loop
```

Gibbs sampling

The Gibbs sampling algorithm builds a Markov chain by iterating through a sequence of conditional distributions.




Figure 3: Sampling trajectory for a bivariate target using Gibbs sampling.

Gibbs sampler

Split the parameter vector $\boldsymbol{\theta} \in \Theta \subseteq \mathbb{R}^p$ into $m \leq p$ blocks,

$$\boldsymbol{\theta}^{[j]} \quad j = 1, \dots, m$$

such that, conditional on the remaining components of the parameter vector $\boldsymbol{\theta}^{-[j]}$, the conditional posterior

$$p(\boldsymbol{\theta}^{[j]} \mid \boldsymbol{\theta}^{-[j]}, \mathbf{y})$$

is from a known distribution from which we can easily simulate.

Gibbs sampling update

At iteration t , we can update each block in turn: note that the k th block uses the partially updated state

$$\boldsymbol{\theta}^{-[k]\star} = (\boldsymbol{\theta}_t^{[1]}, \dots, \boldsymbol{\theta}_t^{[k-1]}, \boldsymbol{\theta}_{t-1}^{[k+1]}, \boldsymbol{\theta}_{t-1}^{[m]})$$

which corresponds to the current value of the parameter vector after the updates.

Notes on Gibbs sampling

- Special case of Metropolis–Hastings with conditional density as proposal q .
- The benefit is that all proposals get accepted, $R = 1$!
- No tuning parameter, but parametrization matters.
- Automatic acceptance does not equal efficiency.

To check the validity of the Gibbs sampler, see the methods proposed in Geweke ([2004](#)).

Efficiency of Gibbs sampling

As the dimension of the parameter space increases, and as the correlation between components becomes larger, the efficiency of the Gibbs sampler degrades

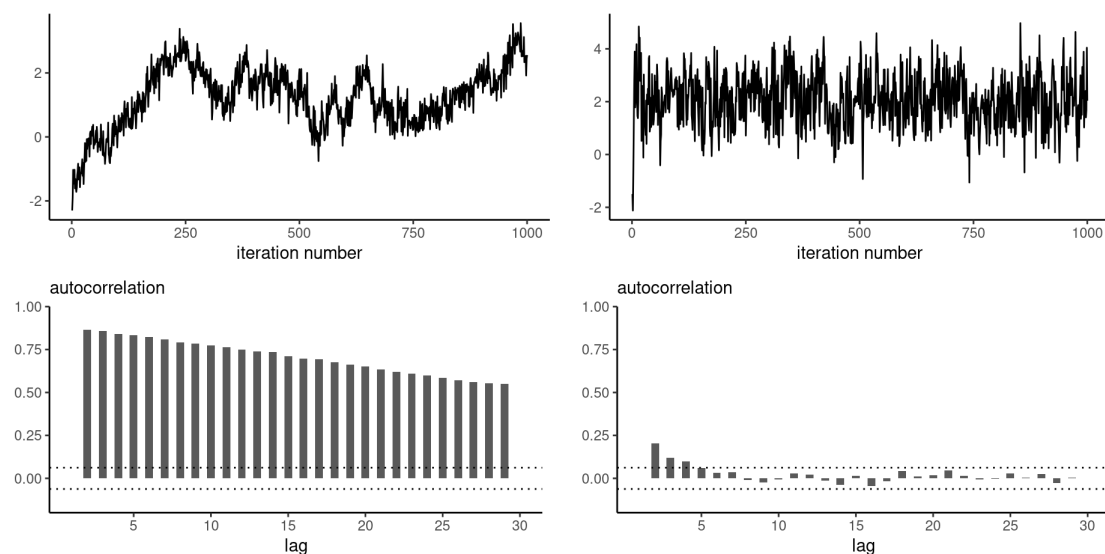


Figure 4: Trace plots (top) and correlograms (bottom) for the first component of a Gibbs sampler with $d = 20$ equicorrelated Gaussian variates with correlation $\rho = 0.9$ (left) and $d = 3$ with equicorrelation $\rho = 0.5$ (right).

Gibbs sampling requires work!

- You need to determine all of the relevant conditional distributions, which often relies on setting conditionally conjugate priors.
- In large models with multiple layers, full conditionals may only depend on a handful of parameters (via directed acyclic graph and moral graph of the model; not covered).

Example of Gibbs sampling

Consider independent and identically distributed observations, with

$$Y_i \sim \text{Gauss}(\mu, \tau), \quad i = 1, \dots, n$$

$$\mu \sim \text{Gauss}(\nu, \omega)$$

$$\tau \sim \text{inv. gamma}(\alpha, \beta)$$

The joint posterior is not available in closed form, but the independent priors for the mean and variance of the observations are conditionally conjugate.

Joint posterior for Gibbs sample

Write the posterior density as usual,

$$\begin{aligned} p(\mu, \tau \mid \mathbf{y}) &\propto \tau^{-\alpha-1} \exp(-\beta/\tau) \\ &\times \tau^{-n/2} \exp \left\{ -\frac{1}{2\tau} \left(\sum_{i=1}^n y_i^2 - 2\mu \sum_{i=1}^n y_i + n\mu^2 \right) \right\} \\ &\times \exp \left\{ -\frac{(\mu - \nu)^2}{2\omega} \right\} \end{aligned}$$

Recognizing distributions from posterior

Consider the conditional densities of each parameter in turn (up to proportionality):

$$p(\mu \mid \tau, \mathbf{y}) \propto \exp \left\{ -\frac{1}{2} \left(\frac{\mu^2 - 2\mu\bar{y}}{\tau/n} + \frac{\mu^2 - 2\nu\mu}{\omega} \right) \right\}$$
$$p(\tau \mid \mu, \mathbf{y}) \propto \tau^{-n/2-\alpha-1} \exp \left[-\frac{\left\{ \frac{\sum_{i=1}^n (y_i - \mu)^2}{2} + \beta \right\}}{\tau} \right]$$

Gibs sample

We can simulate in turn

$$\mu_t \mid \tau_{t-1}, \mathbf{y} \sim \text{Gauss} \left(\frac{n\bar{y}\omega + \tau\nu}{\tau + n\omega}, \frac{\omega\tau}{\tau + n\omega} \right)$$

$$\tau_t \mid \mu_t, \mathbf{y} \sim \text{inv. gamma} \left\{ \frac{n}{2} + \alpha, \frac{\sum_{i=1}^n (y_i - \mu)^2}{2} + \beta \right\}$$

Data augmentation and auxiliary variables

When the likelihood $p(\mathbf{y}; \boldsymbol{\theta})$ is intractable or costly to evaluate (e.g., mixtures, missing data, censoring), auxiliary variables are introduced to simplify calculations.

Consider auxiliary variables $\mathbf{U} \in \mathbb{R}^k$ such that

$$\int_{\mathbb{R}^k} p(\mathbf{U}, \boldsymbol{\theta} \mid \mathbf{y}) d\mathbf{U} = p(\boldsymbol{\theta} \mid \mathbf{y}),$$

i.e., the marginal distribution is that of interest, but evaluation of $p(\mathbf{U}, \boldsymbol{\theta}; \mathbf{y})$ is cheaper.

Bayesian augmentation

The data augmentation algorithm (Tanner & Wong, 1987) consists in running a Markov chain on the augmented state space (Θ, \mathbb{R}^k) , simulating in turn from the conditionals

- $p(\mathbf{U} \mid \boldsymbol{\theta}, \mathbf{y})$ and
- $p(\boldsymbol{\theta} \mid \mathbf{U}, \mathbf{y})$

For more details and examples, see Dyk & Meng (2001) and Hobert (2011).

Data augmentation: probit example

Consider independent binary responses \mathbf{Y}_i , with

$$p_i = \Pr(Y_i = 1) = \Phi(\beta_0 + \beta_1 X_{i1} + \cdots + \beta_p X_{ip}),$$

where Φ is the distribution function of the standard Gaussian distribution. The likelihood of the probit model is

$$L(\boldsymbol{\beta}; \mathbf{y}) = \prod_{i=1}^n p_i^{y_i} (1 - p_i)^{1-y_i},$$

and this prevents easy simulation.

Probit augmentation

We can consider a data augmentation scheme where $Y_i = \mathbf{I}(Z_i > 0)$, where $Z_i \sim \text{Gauss}(\mathbf{x}_i\boldsymbol{\beta}, 1)$, where \mathbf{x}_i is the i th row of the design matrix.

The augmented data likelihood is

$$p(\mathbf{z}, \mathbf{y} \mid \boldsymbol{\beta}) \propto \exp \left\{ -\frac{1}{2} (\mathbf{z} - \mathbf{X}\boldsymbol{\beta})^\top (\mathbf{z} - \mathbf{X}\boldsymbol{\beta}) \right\} \\ \times \prod_{i=1}^n \mathbf{I}(z_i > 0)^{y_i} \mathbf{I}(z_i \leq 0)^{1-y_i}$$

Conditional distributions for probit regression

$$\boldsymbol{\beta} \mid \mathbf{z}, \mathbf{y} \sim \text{Gauss} \left\{ \hat{\boldsymbol{\beta}}, (\mathbf{X}^\top \mathbf{X})^{-1} \right\}$$

$$Z_i \mid y_i, \boldsymbol{\beta} \sim \begin{cases} \text{trunc. Gauss}(\mathbf{x}_i \boldsymbol{\beta}, -\infty, 0) & y_i = 0 \\ \text{trunc. Gauss}(\mathbf{x}_i \boldsymbol{\beta}, 0, \infty) & y_i = 1. \end{cases}$$

with $\hat{\boldsymbol{\beta}} = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{z}$ the ordinary least square estimator.

Data augmentation with scale mixture of Gaussian

The Laplace distribution with mean μ and scale σ has density

$$f(x; \mu, \sigma) = \frac{1}{2\sigma} \exp \left(-\frac{|x - \mu|}{\sigma} \right),$$

and can be expressed as a scale mixture of Gaussians, where $Y \mid \tau \sim \text{Laplace}(\mu, \tau)$ is equivalent to $Z \mid \tau \sim \text{Gauss}(\mu, \tau)$ and $\tau \sim \text{expo}\{(2\sigma)^{-1}\}$.

Joint posterior for Laplace model

With $p(\mu, \sigma) \propto \sigma^{-1}$, the joint posterior for the i.i.d. sample is

$$p(\boldsymbol{\tau}, \mu, \sigma \mid \mathbf{y}) \propto \left(\prod_{i=1}^n \tau_i \right)^{-1/2} \exp \left\{ -\frac{1}{2} \sum_{i=1}^n \frac{(y_i - \mu)^2}{\tau_i} \right\} \\ \times \frac{1}{\sigma^{n+1}} \exp \left(-\frac{1}{2\sigma} \sum_{i=1}^n \tau_i \right)$$

Conditional distributions

The conditionals for $\mu \mid \dots$ and $\sigma \mid \dots$ are, as usual, Gaussian and inverse gamma, respectively. The variances, τ_j , are conditionally independent of one another, with

$$p(\tau_j \mid \mu, \sigma, y_j) \propto \tau_j^{-1/2} \exp \left\{ -\frac{1}{2} \frac{(y_j - \mu)^2}{\tau_j} - \frac{1}{2} \frac{\tau_j}{\sigma} \right\}$$

Inverse transformation

With the change of variable $\xi_j = 1/\tau_j$, we have

$$p(\xi_j \mid \mu, \sigma, y_j) \propto \xi_j^{-3/2} \exp \left\{ -\frac{1}{2\sigma} \frac{\xi_j (y_j - \mu)^2}{\sigma} - \frac{1}{2} \frac{1}{\xi_j} \right\}$$

and we recognize the Wald (or inverse Gaussian) density, where $\xi_i \sim \text{Wald}(\nu_i, \lambda)$ with $\nu_i = \{\sigma / (y_i - \mu)^2\}^{1/2}$ and $\lambda = \sigma^{-1}$.

Bayesian LASSO

Park & Casella (2008) use this hierarchical construction to defined the Bayesian LASSO. With a model matrix \mathbf{X} whose columns are standardized to have mean zero and unit standard deviation, we may write

$$\begin{aligned} \mathbf{Y} \mid \mu, \boldsymbol{\beta}, \sigma^2 &\sim \text{Gauss}_n(\mu \mathbf{1}_n + \mathbf{X}\boldsymbol{\beta}, \sigma \mathbf{I}_n) \\ \beta_j \mid \sigma, \tau &\sim \text{Gauss}(0, \sigma\tau) \\ \tau &\sim \text{expo}(\lambda/2) \end{aligned}$$

Comment about Bayesian LASSO

- If we set an improper prior $p(\mu, \sigma) \propto \sigma^{-1}$, the resulting conditional distributions are all available and thus the model is amenable to Gibbs sampling.
- The Bayesian LASSO places a Laplace penalty on the regression coefficients, with lower values of λ yielding more shrinkage.
- Contrary to the frequentist setting, none of the posterior draws of β are exactly zero.

Summary

- Gibbs sampling is a special case of Metropolis–Hastings algorithm that leads to acceptance
- We need to get the conditional distribution

References

- Dyk, D. A. van, & Meng, X.-L. (2001). The art of data augmentation. *Journal of Computational and Graphical Statistics*, 10(1), 1–50. <https://doi.org/10.1198/10618600152418584>
- Geweke, J. (2004). Getting it right: Joint distribution tests of posterior simulators. *Journal of the American Statistical Association*, 99(467), 799–804. <https://doi.org/10.1198/016214504000001132>
- Hobert, J. P. (2011). The data augmentation algorithm: Theory and methodology. In S. Brooks, A. Gelman, G. Jones, & X. L. Meng (Eds.), *Handbook of Markov chain Monte Carlo* (pp. 253–294). CRC Press. <https://doi.org/10.1201/b10905>
- Park, T., & Casella, G. (2008). The Bayesian Lasso. *Journal of the American Statistical Association*, 103(482), 681–686. <https://doi.org/10.1198/016214508000000337>
- Rue, H., & Held, L. (2005). *Gaussian Markov random fields: Theory and applications* (p. 280). CRC Press.
- Sherlock, C. (2013). Optimal scaling of the random walk Metropolis: General criteria for the 0.234 acceptance rule. *Journal of Applied Probability*, 50(1), 1–15. <https://doi.org/10.1239/jap/1363784420>