

# GRAPHAM: GRAPHICAL MODELS WITH ADAPTIVE RANDOM WALK METROPOLIS ALGORITHMS

MATTI VIHOLA

ABSTRACT. Recently developed adaptive Markov chain Monte Carlo (MCMC) methods have been applied successfully to many problems in Bayesian statistics. Grapham is a new open source implementation covering several such methods, with emphasis on graphical models for directed acyclic graphs. The implemented algorithms include the seminal Adaptive Metropolis algorithm adjusting the proposal covariance according to the history of the chain and a Metropolis algorithm adjusting the proposal scale based on the observed acceptance probability. Different variants of the algorithms allow one, for example, to use these two algorithms together, employ delayed rejection and adjust several parameters of the algorithms. The implemented Metropolis-within-Gibbs update allows arbitrary sampling blocks. The software is written in C and uses a simple extension language Lua in configuration.

## 1. INTRODUCTION

Markov chain Monte Carlo (MCMC) is a general framework for computing expectations over complicated distributions in general state spaces. The methods are based on constructing a Markov chain  $(X_n)_{n \geq 1}$  so that the ergodic averages  $I_N = N^{-1} \sum_{k=1}^N f(X_k)$  converge to  $I = \int f(x)\pi(x)dx$  as  $N \rightarrow \infty$ , where  $\pi$  is the target distribution of interest. Such a chain is often easy to construct using the Metropolis-Hastings algorithm; see, for example, Robert and Casella [1999]. Depending on  $\pi$ , however, it may be difficult to design a practical algorithm so that  $I_N$  would approximate  $I$  well with a moderate number of samples  $N$ .

Recently proposed adaptive MCMC algorithms adjust the parameters of the algorithm (the proposal distribution) on-the-fly, aiming to allow efficient simulation. They have attracted increasing attention in the last few years, after Haario et al. [2001] presented the seminal Adaptive Metropolis (AM) algorithm, and Andrieu and Robert [2001] related adaptive MCMC to the general context of the Robbins-Monro stochastic approximation. After that, several authors have proposed new algorithms and variations, and provided theoretical validation of the methods [Haario et al., 2005, 2006, Atchadé and Rosenthal, 2005, Andrieu and Moulines, 2006, Roberts and Rosenthal, 2009, 2007, Saksman and Vihola, 2008, Atchadé and Fort, 2008, Bai et al., 2008, Vihola, 2009]; see also the recent review by Andrieu and Thoms [2008] and references therein.

Grapham is an open source implementation of several adaptive MCMC algorithms based on the random walk Metropolis sampler. The purpose of Grapham

---

*Date:* September 2, 2009.

*Key words and phrases.* Adaptive Markov chain Monte Carlo, Bayesian statistics, computational method, graphical model, statistical software.

is to provide an experimental tool for evaluating the performance of such algorithms with practical problems, especially in Bayesian statistics. The source code of the software and additional documentation are available for downloading in <http://iki.fi/mvihola/grapham/>.

Rosenthal [2007] describes another adaptive MCMC implementation: AMCMC. It is an R interface to one adaptive MCMC algorithm (referred to as ‘ASCM’ in Section 2 below). Grapham differs from AMCMC in that it relies on a hierarchical model specification and provides more alternative algorithms. Unlike AMCMC, Grapham also provides a set of ready-made standard distribution functions the user can employ as a part of their model specification. This is intended to allow faster development while permitting the user to define arbitrary distributions easily.

The models are specified in Grapham by defining a set of variables with their conditional distributions. Such models are often referred to as ‘graphical models’; see, for example, Lauritzen [1996] and references therein. This underlying philosophy of Grapham reminds that of BUGS [Spiegelhalter et al., 1996–2008]; see also the review Murphy [2007] of other software for graphical models. The advantage of Grapham over BUGS is that the adaptive MCMC algorithms can be much more efficient than the non-adaptive (Metropolis-within-)Gibbs algorithms of BUGS. One should, however, notice that Grapham is an experimental project not offering the versatility and maturity of BUGS. It is also likely that BUGS performs better than Grapham with many simpler models.

## 2. ALGORITHMS

The general form of the algorithms implemented in Grapham can be described as follows. Let  $X_0 \equiv x_0 \in \mathbb{R}^d$  be a given starting point for the state chain, and  $\theta_0$  and  $L_0$  stand for the initial scaling parameter and the (lower-diagonal with non-zero diagonal) shape matrix, respectively. For  $n \geq 1$ , the recursion follows:

- (S1) form a proposal  $Y_n = X_{n-1} + \theta_{n-1}L_{n-1}W_n$ , where  $W_n$  is an independent sample from a symmetric proposal distribution  $q_0$ ,
- (S2) with probability  $\alpha_n = \min\{1, \pi(Y_n)/\pi(X_{n-1})\}$ , the proposal is accepted and  $X_n = Y_n$ ; otherwise, the proposal is rejected and  $X_n = X_{n-1}$ , and
- (S3) update the scaling parameter  $\theta_{n-1} \rightarrow \theta_n > 0$  and the shape  $L_{n-1} \rightarrow L_n \in \mathbb{R}^{d \times d}$  according to the selected adaptive algorithm.

The steps (S1) and (S2) implement an iteration of the random-walk Metropolis algorithm with the proposal distribution  $q_0$  scaled by the factor  $\theta_{n-1}L_{n-1}$ . Step (S3) implements the adaptation, changing the scaling parameters  $\theta_n$  and  $L_n$  based on the history of the chain. Examples of such updates are given below.

Instead of applying the iteration (S1)–(S3) at once to all the elements of the vector  $X_n$ , one may use Metropolis-within-Gibbs and apply the iteration sequentially to subsets of the elements of  $X_n$ , as in the single component AM algorithm suggested by Haario et al. [2005]. These sampling blocks can be selected freely in Grapham. The proposal distribution  $q_0$  in (S1) can also be chosen. Grapham currently implements (multivariate) Gaussian, student, uniform (in a cube) and (a  $d$ -fold product of) Laplace proposal distributions.

The adaptation of (S3) depends on the selected algorithm. The Adaptive Metropolis (AM) algorithm of Haario et al. [2001] implies constant scaling  $\theta_n = \theta_0$  for all  $n \geq 1$ . The shape matrix  $L_n$  is the Cholesky factor of a covariance estimate of the chain. In particular,  $L_n L_n^T = C_n$  with a positive definite  $C_0 \in \mathbb{R}^{d \times d}$  and defined through

$$\begin{aligned} (1) \quad M_n &= (1 - \eta_n)M_{n-1} + \eta_n X_n \quad \text{and} \\ (2) \quad C_n &= (1 - \eta_n)C_{n-1} + \eta_n (X_n - M_{n-1})(X_n - M_{n-1})^T, \end{aligned}$$

with  $M_0 \equiv x_0$ . The weight sequence  $\eta_n \in (0, 1)$  can be selected arbitrarily, but it is recommended to choose  $\eta_n$  decaying to zero. For example, setting  $\eta_n = \eta_0 > 0$  for all  $n \geq 1$  results in an algorithm similar to the Adaptive Proposal (AP) algorithm [Haario et al., 1999]. This algorithm does not, in general, provide valid simulation; see the example in Haario et al. [2001]. The original AM algorithm employs the default value  $\eta_n = (n + 1)^{-1}$ , in which case  $M_n$  and  $C_n$  coincide with the average and (asymptotically) the sample covariance of  $X_0, \dots, X_n$ , respectively. The updated Cholesky factor  $L_{n+1}$  of  $C_{n+1}$  is computed efficiently from  $L_n$  by a rank one update requiring  $O(d^2)$  operations [Dongarra et al., 1979]. Observe that the same order of operations is needed when forming the proposal  $Y_n$  in (S1).

The adaptive scaling Metropolis (ASCM) algorithm as proposed by Atchadé and Rosenthal [2005] and Roberts and Rosenthal [2007, 2009] leaves the shape matrix constant  $L_n = L_0$  for all  $n \geq 1$ . The scaling parameter  $\theta_n$  is updated according to the observed acceptance probability. The default update in Grapham is

$$(3) \quad \theta_n = \theta_{n-1} \left[ 1 + \eta_n \left( \frac{\alpha_n}{\alpha^*} - 1 \right) \right],$$

where  $\alpha^*$  is the desired acceptance probability. The default values for  $\alpha^*$  are 0.44 in dimension one and 0.234 otherwise following Roberts and Rosenthal [2009]. The user can also supply an alternative, arbitrary update function easily, as exemplified in Section 4.

These two algorithms, AM and ASCM, can be used simultaneously, as suggested in Atchadé and Fort [2008] and Andrieu and Thoms [2008]. Additional flavours to the algorithms include a Rao-Blackwellised version of AM [Andrieu and Thoms, 2008] modifying the update formulae (1) and (2) to

$$\begin{aligned} (4) \quad M_n &= (1 - \eta_n)M_{n-1} + \eta_n [\alpha_n Y_n + (1 - \alpha_n)X_{n-1}] \quad \text{and} \\ (5) \quad C_n &= (1 - \eta_n)C_{n-1} + \eta_n [\alpha_n (Y_n - M_{n-1})(Y_n - M_{n-1})^T \\ &\quad + (1 - \alpha_n)(X_{n-1} - M_{n-1})(X_{n-1} - M_{n-1})^T]. \end{aligned}$$

There is a possibility to use (two-stage) delayed rejection (DR) with AM [Haario et al., 2006]. DR can also be applied when using ASCM, so that only the first-stage acceptance probability  $\alpha_n$  is employed in (3).

Grapham implements three different burn-in strategies for adaptation. The default ‘greedy’ strategy performs continuous adaptation during the whole MCMC run. The ‘traditional’ strategy as proposed in Haario et al. [2001] uses a fixed proposal for the burn-in and then performs continuous adaptation during the rest of the simulation. One may also apply a ‘freeze’ strategy adapting only during the burn-in and keeping the obtained parameters fixed during the estimation run.

It is possible to employ a mixture of two proposal density components, a fixed and an adaptive one [Roberts and Rosenthal, 2009]. This is implemented in Grapham so that, with probability  $p_{\text{mix}}$ , the initial parameters  $L_0$  and  $\theta_0$  are used in (S1) instead of the adapted values  $\theta_{n-1}$  and  $L_{n-1}$ . The user may define also a non-constant mixing probability  $p_{\text{mix}}^{(n)} \in [0, 1]$ . This feature can be used, for example, to introduce a ‘gradual burn-in,’ by defining a decaying sequence  $p_{\text{mix}}^{(n)} \rightarrow 0$ .

### 3. IMPLEMENTATION

Grapham does not have an interactive ‘user interface.’ It is simply executed from the command prompt (shell) with input file names as parameters. The input files contain the model specification and the simulation parameters. It is also possible to define the functional of interest in the input files. For more complicated functionals, however, it may be convenient to store (a subset of) the samples simulated by Grapham and process them in another environment. The samples can be saved into a file in the CSV (comma separated values) format or in a simple binary format. The former allows the samples to be easily imported to many other environments. There are ready-made functions for loading the binary data files into R [R Development Core Team, 2009], Matlab<sup>®</sup> (The MathWorks, Natick, Massachusetts) and Octave [Eaton, 2002] environments.

The core of Grapham is implemented in C. It includes some numerical Fortran subroutines from the Netlib repository [Browne et al., 1995] and can optionally be compiled with the dSFMT random number generator of Saito and Matsumoto [2008] instead of using the random number generators provided by the C standard libraries. The configuration of Grapham is done using the small and publicly available extension language Lua [Ierusalimsky et al., 1996]. While minimalistic, Lua is in fact a full-featured programming language offering a great flexibility. For example, the user can supply functions as configuration parameters and apply data from external files in the model. In fact, Grapham includes some functions written in Lua, for example for reading data files in the CSV format. The Numeric Lua package [Carvalho, 2005] can also be compiled with Grapham to allow easy working with vector-valued variables.

There are numerous ready-made distribution functions available for defining the conditional densities associated with the variables. The densities can also be defined arbitrarily as Lua functions. Likewise, the functional of interest may be written in Lua. However, to allow optimal performance, Grapham allows the user to supply densities and functionals in a separate C library with ease.

### 4. AN EXAMPLE SESSION

Consider the baseball model of Rosenthal [1996] used as an example also with AMCMC [Rosenthal, 2007]. It consists of 38 real-valued variables, defined hierarchically as depicted in Fig. 1. The file specifying this model in Grapham is shown in Fig. 2. The model is defined in the Lua table `model`, defined in lines 4–20. Each variable is defined by an entry containing a logarithmic density, conditional on the parent variables. The variables  $\mu$ ,  $t$  and  $y$  in the example have standard distributions:  $\mu$  has (an improper) uniform distribution over  $\mathbb{R}$ , while




FIGURE 1. The graphical representation of the baseball model. The nodes with observed values ('data') are shown in grey.

```

1  const = {
2    v = 0.00434
3  }
4  model = {
5    mu = {
6      density = "duniform"
7    },
8    t = {
9      parents = {"mu","a"}, density = "dnorm"
10   },
11   y = {
12     parents = {"t", "v"}, density = "dnorm"
13   },
14   a = {
15     init_val = 1,
16     density = function(a_)
17       return dexp(1/a_, 1/2)
18     end
19   },
20 }
21 _, y = read_csv("models/baseball.data")
22 repeat_block({"y","t"}, y[1])
23 function functional()
24   return {t1, mu, a}
25 end
26 para = {
27   niter = 30000, nburn = 10000, algorithm = "ascm",
28 }

```

FIGURE 2. The Lua code in the file models/baseball.lua specifying the model of Fig. 1 in Grapham.

$t$  and  $y$  are conditionally Gaussian with means  $\mu$  and  $t$  and variances  $a$  and  $v$ , respectively. The reciprocal of the variable  $a$  is exponentially distributed; this is defined through a Lua function defined in lines 16–18, calling `dexp`, the exponential distribution function. The model is, in fact, then modified by the function `repeat_block`. The block of variables  $(y, t)$  in the model is replicated 18 times to

```

1 para.scaling_adapt = function(sc, alpha, dim, k)
2   if alpha>0.44 then
3     delta = 1
4   else
5     delta = -1
6   end
7   return sc*exp(delta*min(0.01, 1/sqrt(k+1)))
8 end
9 para.dr = 0.1; para.proposal = "student"

```

FIGURE 3. The Lua code in the file `models/amcmc_dr.lua`.

obtain blocks  $(y_1, t_1), \dots, (y_{18}, t_{18})$ . At the same time, the function `repeat_block` sets the values of  $y_i$  to the 18 values read from the CSV file `baseball.data` using the function `read_csv`.

The following shows an example run of Grapham with the model specification of Fig. 2.

```

$ ./grapham models/baseball.lua
Functional average = [ 0.392507 0.267393 0.318917 ]
Acceptance rates: ( a ): 44.03% ( t7 ): 43.97% ( t9 ): 44.00%

```

The part of the output shown above contains the computed estimate of the expected value of the functional specified in lines 23–25 of Fig. 2, that is, the mean of the vector  $[t_1, \mu, a]$ , giving a similar estimate for  $t_1$  as obtained by AMCMC [Rosenthal, 2007]. Moreover, the average acceptance rate of the nodes was approximately 44%, which is the default value of the desired acceptance probability  $\alpha^*$ . The run consisted of 40000 (of which 10000 burn-in) iterations using the ASCM algorithm for each real-valued variable at a time. This algorithm is very similar to the one implemented in AMCMC. The running time of Grapham was approximately 1.0 seconds with Intel Pentium 4 at 2.80GHz. As a comparison, the same run with AMCMC (with both the density and the functional specified in C for optimal performance) took approximately 3.8 seconds. The faster simulation speed of Grapham is explained by the hierarchical model setup, which Grapham can take advantage of. That is, only part of the conditional densities in the target distribution need to be evaluated when each variable is updated.

Let us modify the above example, by adding the lines shown in Fig. 3 to the model specification of Fig. 2. The supplied function `para.scaling_adapt` replaces the default update in (3), and in fact implements exactly the scaling adaptation algorithm of AMCMC [Rosenthal, 2007, Roberts and Rosenthal, 2009]. The value set to the parameter `para.dr` means that delayed rejection is used, with a 0.1 times down-scaled proposal in the second stage. Moreover, instead of the default Gaussian distribution, the proposal samples are drawn from a Student's  $t$ -distribution.

```

$ ./grapham models/baseball.lua models/amcmc_dr.lua
Functional average = [ 0.392465 0.266204 0.321466 ]
Acceptance rates: ( a ): 70.26% (47.49%/22.77%) ( t7 ): 70.66%
(47.55%/23.10%) ( t9 ): 71.08% (47.62%/23.46%)

```

In this case, the total acceptance rate of each block is around 70%, of which roughly two thirds are accepted in the first stage and one third in the second, delayed rejection stage. The estimates obtained for  $t_1$ ,  $\mu$  and  $a$  appear similar to the first run.

Finally, to exemplify how the data simulated by Grapham can be used in other environments, let us run Grapham with the command line

```
$ ./grapham models/baseball.lua -e "para.outfile='bb.bin'"
```

This command includes the chunk of Lua code `para.outfile='bb.bin'` after reading the file `baseball.lua`. As a consequence, the simulated samples are written in the file `bb.bin`. In R, one could, for example, write

```
> source("tools/grapham_read.r")
> data <- grapham_read("bb.bin", nthin=10)
> plot(data$a, data$t1)
```

which would plot every tenth of the 30000 simulated samples of  $(a, t_1)$  in the same figure.

## 5. CONCLUSIONS

Grapham provides a flexible open-source test bed for evaluating the performance of different adaptive random walk Metropolis algorithms, especially with hierarchical models often encountered in Bayesian statistics. It provides a fairly simple and general way of determining models and functionals and for incorporating data into the model. The simulation speed of Grapham is good, even in a relatively high-dimensional setting, as the implemented algorithms involve at most a quadratic number of operations with respect to the dimension. The user has extensive control over the various parameters of the algorithms, enabling a thorough testing of different adaptation strategies. Moreover, new adaptive algorithms of the similar random walk type can be easily added to Grapham.

## ACKNOWLEDGEMENTS

The author thanks Professor Antti Penttinen and the referees for helpful comments on the manuscript. This work was supported by the Academy of Finland, projects no. 110599 and 201392, by the Finnish Academy of Science and Letters, the Vilho, Yrjö and Kalle Väisälä Foundation, by the Finnish Centre of Excellence in Analysis and Dynamics Research and by the Finnish Graduate School in Stochastics and Statistics.

## REFERENCES

- C. Andrieu and É. Moulines. On the ergodicity properties of some adaptive MCMC algorithms. *Ann. Appl. Probab.*, 16(3):1462–1505, 2006.
- C. Andrieu and C. P. Robert. Controlled MCMC for optimal sampling. Technical Report Ceremade 0125, Université Paris Dauphine, 2001.
- C. Andrieu and J. Thoms. A tutorial on adaptive MCMC. *Statist. Comput.*, 18(4):343–373, Dec. 2008.
- Y. Atchadé and G. Fort. Limit theorems for some adaptive MCMC algorithms with subgeometric kernels. URL <http://arxiv.org/abs/0807.2952>. Preprint, July 2008.

- Y. F. Atchadé and J. S. Rosenthal. On adaptive Markov chain Monte Carlo algorithms. *Bernoulli*, 11(5):815–828, 2005.
- Y. Bai, G. O. Roberts, and J. S. Rosenthal. On the containment condition for adaptive Markov chain Monte Carlo algorithms. URL <http://probability.ca/jeff/research.html>. Preprint, July 2008.
- S. Browne, J. Dongarra, E. Grosse, and T. Rowan. The Netlib mathematical software repository. Technical report, Corporation for National Research Initiatives, 1995. URL <http://netlib.org/>.
- L. Carvalho. Numeric Lua project page, 2005. URL <http://luaforge.net/projects/numlua/>. Visited 11th June 2009.
- J. J. Dongarra, J. R. Bunch, C. B. Moler, and G. W. Stewart. *LINPACK Users' Guide*. Society for Industrial and Applied Mathematics, 1979. ISBN 0-89871-172-X.
- J. W. Eaton. *GNU Octave Manual*. Network Theory Limited, 2002. ISBN 0-9541617-2-6. URL <http://www.gnu.org/software/octave/>.
- H. Haario, E. Saksman, and J. Tamminen. Adaptive proposal distribution for random walk Metropolis algorithm. *Comput. Statist.*, 14(3):375–395, 1999.
- H. Haario, E. Saksman, and J. Tamminen. An adaptive Metropolis algorithm. *Bernoulli*, 7(2):223–242, 2001.
- H. Haario, E. Saksman, and J. Tamminen. Componentwise adaptation for high dimensional MCMC. *Comput. Statist.*, 20(2):265–274, 2005.
- H. Haario, M. Laine, A. Mira, and E. Saksman. DRAM: Efficient adaptive MCMC. *Statist. Comput.*, 16(4):339–354, Dec. 2006.
- R. Ierusalimschy, L. H. de Figueiredo, and W. C. Filho. Lua—an extensible extension language. *Software: Practice and Experience*, 26(6):635–652, June 1996. URL <http://www.lua.org/>.
- S. L. Lauritzen. *Graphical Models*. Oxford University Press, 1996. ISBN 0-19-852219-3.
- K. Murphy. Software for graphical models: A review. *ISBA Bulletin*, 14(4), Dec. 2007.
- R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2009. URL <http://www.R-project.org>.
- C. P. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer-Verlag, New York, 1999. ISBN 0-387-98707-X.
- G. O. Roberts and J. S. Rosenthal. Coupling and ergodicity of adaptive Markov chain Monte Carlo algorithms. *J. Appl. Probab.*, 44(2):458–475, 2007.
- G. O. Roberts and J. S. Rosenthal. Examples of adaptive MCMC. *J. Comput. Graph. Statist.*, 18(2):349–367, 2009.
- J. S. Rosenthal. AMCMC: An R interface for adaptive MCMC. *Comput. Statist. Data Anal.*, 51(12), Aug. 2007. URL <http://www.probability.ca/amcmc/>.
- J. S. Rosenthal. Analysis of the Gibbs sampler for a model related to James-Stein estimators. *Statist. Comput.*, 6(3):269–275, 1996.
- M. Saito and M. Matsumoto. SIMD-oriented fast Mersenne twister: a 128-bit pseudorandom number generator. In *Monte Carlo and Quasi-Monte Carlo Methods 2006*, pages 607–622. Springer, 2008.



- E. Saksman and M. Vihola. On the ergodicity of the adaptive Metropolis algorithm on unbounded domains. URL <http://arxiv.org/abs/0806.2933>. Preprint, June 2008.
- D. J. Spiegelhalter, A. Thomas, N. G. Best, W. R. Gilks, and D. Lunn. BUGS: Bayesian inference using Gibbs sampling, 1996–2008. URL <http://www.mrc-bsu.cam.ac.uk/bugs/>.
- M. Vihola. On the stability and ergodicity of an adaptive scaling Metropolis algorithm. URL <http://arxiv.org/abs/0903.4061>. Preprint, Mar. 2009.

MATTI VIHOLA, DEPARTMENT OF MATHEMATICS AND STATISTICS, P.O. BOX 35 (MAD),  
FI-40014 UNIVERSITY OF JYVÄSKYLÄ, FINLAND

*E-mail address:* [matti.vihola@iki.fi](mailto:matti.vihola@iki.fi)

*URL:* <http://iki.fi/mvihola/>