

Package ‘tmg’

February 20, 2015

Type Package

Title Truncated Multivariate Gaussian Sampling

Version 0.3

Date 2015-02-11

Author Ari Pakman

Maintainer Ari Pakman <ari@stat.columbia.edu>

Description Random number generation of truncated multivariate Gaussian distributions using Hamiltonian Monte Carlo. The truncation is defined using linear and/or quadratic polynomials.

License GPL-2

Imports Rcpp

LinkingTo Rcpp, RcppEigen

URL <http://arxiv.org/abs/1208.4118>

NeedsCompilation yes

Repository CRAN

Date/Publication 2015-02-11 20:57:16

R topics documented:

rtmg	1
Index	4

rtmg *Sample from truncated multivariate Gaussians*

Description

This function generates samples from a Markov chain whose equilibrium distribution is a d-dimensional multivariate Gaussian truncated by linear and quadratic inequalities. The probability log density is

$$\log p(X) = -0.5 X^T M X + r^T X + \text{const}$$

in terms of a precision matrix M and a vector r . The constraints are imposed as explained below. The Markov chain is built using the Hamiltonian Monte Carlo technique. See the reference below for an explanation of the method.

Usage

```
rtmg(n, M, r, initial, f = NULL, g = NULL, q = NULL, burn.in = 30)
```

Arguments

<code>n</code>	Number of samples.
<code>M</code>	A d-by-d precision matrix of the multivariate Gaussian density.
<code>r</code>	A d-dimensional vector for the linear coefficient of the log density.
<code>initial</code>	A d-dimensional vector with the initial value of the Markov chain. Must satisfy the truncation inequalities strictly.
<code>f</code>	An m-by-d matrix, where m is the number of linear constraints. The constraints require each component of the m-dimensional vector $fX + g$ to be non-negative.
<code>g</code>	An m-dimensional vector with the constant terms in the above linear constraints.
<code>q</code>	A list of quadratic constraints. Each element i of q is a 3-element list of the form $q[[i]] = \text{list}(A,B,C)$ and imposes a quadratic constraint of the form $X^T A X + B^T X + C \geq 0$ where A is a d-by-d matrix, B is a d-dimensional vector and C is a number.
<code>burn.in</code>	The number of burn-in iterations. The Markov chain is sampled $n + \text{burn.in}$ times, and the last n samples are returned.

Details

For linear constraints, both f and g must be provided. If no values for f , g and q are provided, the function returns samples from an untruncated Gaussian. When the truncation equations define several disconnected regions, the samples belong to the region of the initial value.

Value

An n-by-d matrix, where each row contains a sample.

Author(s)

Ari Pakman

Maintainer: Ari Pakman <ari@stat.columbia.edu>

References

Pakman, A. and Paninski, L., Exact Hamiltonian Monte Carlo for Truncated Multivariate Gaussians
- Journal of Computational and Graphical Statistics, 2014

<http://arxiv.org/abs/1208.4118>

Examples

```
# Set number of samples
n=15000;

#Define precision matrix and linear term
M = matrix(c(.5,-.4, -.4,.5), 2,2)
r = c(0,0)

# Set initial point for the Markov chain
initial = c(4,1)

# Define two linear constraints
f = diag(2)
f[1,2] = 1
g = c(0,0)

# Define two quadratic constraints
A1 = matrix(c(-1/8,0,0,-1/2),2,2)
B1 = c(.5,.5)
C1 = 3/4
constr1 = list(A1,B1,C1)

A2 = matrix(c(4,-1,-1,8),2,2)
B2 = c(0,5)
C2 = -1
constr2 = list(A2,B2,C2)

q = list(constr1,constr2)

# Sample and plot
samples = rtmg(n, M, r, initial, f,g, q);
plot(samples, pch=".")
```

Index

*Topic **distribution**
rtmg, 1

rtmg, 1