

Package ‘elhmc’

October 30, 2024

Type Package

Title Sampling from a Empirical Likelihood Bayesian Posterior of Parameters Using Hamiltonian Monte Carlo

Version 1.2.1

Date 2024-10-29

Description A tool to draw samples from a Empirical Likelihood Bayesian posterior of parameters using Hamiltonian Monte Carlo.

Imports emplik, plyr, stats, MASS, graphics

License GPL-2

RoxygenNote 7.3.2

NeedsCompilation no

Author Dang Trung Kien [aut],
Sanjay Chaudhuri [aut, cre],
Neo Han Wei, [aut]

Maintainer Sanjay Chaudhuri <schaudhuri2@unl.edu>

Repository CRAN

Date/Publication 2024-10-30 00:40:02 UTC

Contents

ELHMC	1
Index	5

ELHMC	<i>Empirical Likelihood Hamiltonian Monte Carlo Sampling</i>
-------	--

Description

This function draws samples from a Empirical Likelihood Bayesian posterior distribution of parameters using Hamiltonian Monte Carlo.

Usage

```

ELHMC(
  initial,
  data,
  fun,
  dfun,
  prior,
  dprior,
  n.samples = 100,
  lf.steps = 10,
  epsilon = 0.05,
  p.variance = 1,
  tol = 10^-5,
  detailed = FALSE,
  print.interval = 1000,
  plot.interval = 0,
  which.plot = NULL,
  FUN,
  DFUN
)

```

Arguments

<code>initial</code>	a vector containing the initial values of the parameters
<code>data</code>	a matrix containing the data
<code>fun</code>	the estimating function g . It takes in a parameter vector <code>params</code> as the first argument and a data point vector <code>x</code> as the second parameter. This function returns a vector.
<code>dfun</code>	a function that calculates the gradient of the estimating function g . It takes in a parameter vector <code>params</code> as the first argument and a data point vector <code>x</code> as the second argument. This function returns a matrix.
<code>prior</code>	a function with one argument <code>x</code> that returns the log joint prior density of the parameters of interest
<code>dprior</code>	a function with one argument <code>x</code> that returns the gradients of the log densities of the parameters of interest
<code>n.samples</code>	number of samples to draw
<code>lf.steps</code>	number of leap frog steps in each Hamiltonian Monte Carlo update
<code>epsilon</code>	the leap frog step size(s). This has to be a single numeric value or a vector of the same length as <code>initial</code> .
<code>p.variance</code>	the covariance matrix of a multivariate normal distribution used to generate the initial values of momentum p in Hamiltonian Monte Carlo. This can also be a single numeric value or a vector. See Details.
<code>tol</code>	EL tolerance
<code>detailed</code>	If this is set to TRUE, the function will return a list with extra information.

<code>print.interval</code>	the frequency at which the results would be printed on the terminal. Defaults to 1000.
<code>plot.interval</code>	the frequency at which the drawn samples would be plotted. The last half of the samples drawn are plotted after each <code>plot.interval</code> steps. The acceptance rate is also plotted. Defaults to 0, which means no plot.
<code>which.plot</code>	the vector of parameters to be plotted after each <code>plot.interval</code> . Defaults to NULL, which means no plot.
<code>FUN</code>	the same as <code>fun</code> but takes in a matrix <code>X</code> instead of a vector <code>x</code> and returns a matrix so that <code>FUN(params, X)[i,]</code> is the same as <code>fun(params, X[i,])</code> . Only one of <code>FUN</code> and <code>fun</code> should be provided. If both are then <code>fun</code> is ignored.
<code>DFUN</code>	the same as <code>dfun</code> but takes in a matrix <code>X</code> instead of a vector <code>x</code> and returns an array so that <code>DFUN(params, X)[, i]</code> is the same as <code>dfun(params, X[i,])</code> . Only one of <code>DFUN</code> and <code>dfun</code> should be provided. If both are then <code>dfun</code> is ignored.

Details

Suppose there are data $x = (x_1, x_2, \dots, x_n)$ where x_i takes values in R^p and follow probability distribution F . Also, F comes from a family of distributions that depends on a parameter $\theta = (\theta_1, \dots, \theta_d)$ and there is a smooth function $g(x_i, \theta) = (g_1(x_i, \theta), \dots, g_q(x_i, \theta))^T$ that satisfies $E_F[g(x_i, \theta)] = 0$ for $i = 1, \dots, n$.

ELHMC draws samples from a Empirical Likelihood Bayesian posterior distribution of the parameter θ , given the data x as `data`, the smoothing function g as `fun`, and the gradient of g as `dfun` or $G(X) = (g(x_1), g(x_2), \dots, g(x_n))^T$ as `FUN` and the gradient of G as `DFUN`.

Value

The function returns a list with the following elements:

<code>samples</code>	A matrix containing the parameter samples
<code>acceptance.rate</code>	The acceptance rate
<code>call</code>	The matched call

If `detailed = TRUE`, the list contains these extra elements:

<code>proposed</code>	A matrix containing the proposed values at <code>n.samples - 1</code> Hamiltonian Monte Carlo updates
<code>acceptance</code>	A vector of TRUE/FALSE values indicates whether each proposed value is accepted
<code>trajectory</code>	A list with 2 elements <code>trajectory.q</code> and <code>trajectory.p</code> . These are lists of matrices containing position and momentum values along trajectory in each Hamiltonian Monte Carlo update.

References

Chaudhuri, S., Mondal, D. and Yin, T. (2017) Hamiltonian Monte Carlo sampling in Bayesian empirical likelihood computation. *Journal of the Royal Statistical Society: Series B*.

Neal, R. (2011) MCMC for using Hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo* (eds S. Brooks, A. Gelman, G. L. Jones and X.-L. Meng), pp. 113-162. New York: Taylor and Francis.

Examples

```
## Suppose there are four data points (1, 1), (1, -1), (-1, -1), (-1, 1)
x = rbind(c(1, 1), c(1, -1), c(-1, -1), c(-1, 1))
## If the parameter of interest is the mean, the smoothing function and
## its gradient would be
f <- function(params, x) {
  x - params
}
df <- function(params, x) {
  rbind(c(-1, 0), c(0, -1))
}
## Draw 50 samples from the Empirical Likelihood Bayesian posterior distribution
## of the mean, using initial values (0.96, 0.97) and standard normal distributions
## as priors:
normal_prior <- function(x) {
  -0.5 * (x[1] ^ 2 + x[2] ^ 2) - log(2 * pi)
}
normal_prior_log_gradient <- function(x) {
  -x
}
set.seed(1234)
mean.samples <- ELHMC(initial = c(0.96, 0.97), data = x, fun = f, dfun = df,
  n.samples = 50, prior = normal_prior,
  dprior = normal_prior_log_gradient)
plot(mean.samples$samples, type = "l", xlab = "", ylab = "")
```

Index

ELHMC, 1