

Package ‘sns’

November 2, 2022

Type Package

Title Stochastic Newton Sampler (SNS)

Version 1.2.2

Date 2022-11-01

Author Alireza S. Mahani, Asad Hasan, Marshall Jiang, Mansour T.A. Sharabiani

Maintainer Alireza Mahani <alireza.s.mahani@gmail.com>

Description Stochastic Newton Sampler (SNS) is a Metropolis-Hastings-based, Markov Chain Monte Carlo sampler for twice differentiable, log-concave probability density functions (PDFs) where the proposal density function is a multivariate Gaussian resulting from a second-order Taylor-series expansion of log-density around the current point. The mean of the Gaussian proposal is the full Newton-Raphson step from the current point. A Boolean flag allows for switching from SNS to Newton-Raphson optimization (by choosing the mean of proposal function as next point). This can be used during burn-in to get close to the mode of the PDF (which is unique due to concavity). For high-dimensional densities, mixing can be improved via 'state space partitioning' strategy, in which SNS is applied to disjoint subsets of state space, wrapped in a Gibbs cycle. Numerical differentiation is available when analytical expressions for gradient and Hessian are not available. Facilities for validation and numerical differentiation of log-density are provided. Note: Formerly available versions of the MfUSampler can be obtained from the archive <<https://cran.r-project.org/src/contrib/Archive/MfUSampler/>>.

License GPL (>= 2)

Imports mvtnorm, coda, numDeriv

Suggests RegressionFactory, MfUSampler

NeedsCompilation no

Repository CRAN

Date/Publication 2022-11-02 11:02:22 UTC

R topics documented:

ess	2
plot.sns	3
predict.sns	4

sns	6
sns.check.logdensity	9
sns.fghEval.numaug	13
sns.make.part	14
sns.run	15
summary.sns	18
Index	20

 ess

Effective Sample Size Calculator

Description

Computes the effective sample size of MCMC chains, using the algorithm in Section 2.3 of the paper by Madeline Thompson. The algorithm is taken from earlier work on ‘Initial Sequence Estimators’ by multiple authors.

Usage

```
ess(x, method = c("coda", "ise"))
```

Arguments

x	Matrix object with each sample (possibly multivariate) as a row. Effective sample size calculation is done independently for each column of x.
method	Method of calculating effective size. Current options are "coda" which calls effectiveSize function in coda package, and "ise" which uses the ‘Initial Sequence Estimators’ method described in Section 2.3 of Thompson (2010).

Value

Vector with effective sample sizes for the time series in each column of x.

Author(s)

Alireza S. Mahani, Asad Hasan, Marshall Jiang, Mansour T.A. Sharabiani

References

Thompson, Madeleine (2010) *A Comparison of Methods for Computing Autocorrelation Time*
<https://arxiv.org/pdf/1011.0175v1.pdf>

`plot.sns`*Plotting "sns" Objects*

Description

Method for visualizing the output of `sns.run`.

Usage

```
## S3 method for class 'sns'  
plot(x, nburnin = max(nrow(x)/2, attr(x, "nnr"))  
     , select = if (length(x) <= 10) 1:5 else 1, ...)
```

Arguments

<code>x</code>	Object of class "sns", typically the output of <code>sns.run</code> .
<code>nburnin</code>	Number of burn-in iterations to discard before generating effective sample size, histograms, and autocorrelation plots.
<code>select</code>	Which plot types must be generated. See below for description.
<code>...</code>	Arguments passed to/from other functions.

Value

`plot.sns` produces the following types of plots: 1) log-probability trace plot (vertical line, if present, indicates transition from nr to mcmc mode), 2) trace plot of state variables (one per coordinate; vertical line has same meaning as 1), 3) effective sample size by coordinate (horizontal line indicates maximum effective size possible, equal to number of samples after discarding `nburnin` initial iterations), 4) post-burnin state vector histograms (one per coordinate, vertical line indicates post-burnin average), 5) autocorrelation plots, one per coordinate.

Author(s)

Alireza S. Mahani, Asad Hasan, Marshall Jiang, Mansour T.A. Sharabiani

References

Mahani A.S., Hasan A., Jiang M. & Sharabiani M.T.A. (2016). Stochastic Newton Sampler: The R Package `sns`. *Journal of Statistical Software, Code Snippets*, 74(2), 1-33. doi:10.18637/jss.v074.c02

See Also

[sns.run](#)

 predict.sns

Sample-based prediction using "sns" Objects

Description

Method for sample-based prediction using the output of [sns.run](#).

Usage

```
## S3 method for class 'sns'
predict(object, fpred
  , nburnin = max(nrow(object)/2, attr(object, "nnr"))
  , end = nrow(object), thin = 1, ...)
## S3 method for class 'predict.sns'
summary(object
  , quantiles = c(0.025, 0.5, 0.975)
  , ess.method = c("coda", "ise"), ...)
## S3 method for class 'summary.predict.sns'
print(x, ...)
```

Arguments

object	Object of class "sns" (output of sns.run) or "predict.sns" (output of <code>predict.sns</code>).
fpred	Prediction function, accepting a single value for the state vector and producing a vector of outputs.
nburnin	Number of burn-in iterations discarded for sample-based prediction.
end	Last iteration used in sample-based prediction.
thin	One out of thin iterations within the specified range are used for sample-based prediction.
quantiles	Values for which sample-based quantiles are calculated.
ess.method	Method used for calculating effective sample size. Default is to call <code>effectiveSize</code> from package <code>coda</code> .
x	An object of class "summary.predict.sns".
...	Arguments passed to/from other functions.

Value

`predict.sns` produces a matrix with number of rows equal to the length of prediction vector produces by `fpred`. Its number of columns is equal to the number of samples used within the user-specified range, and after thinning (if any). `summary.predict.sns` produces sample-based prediction mean, standard deviation, quantiles, and effective sample size.

Note

See package vignette for more details on SNS theory, software, examples, and performance.

Author(s)

Alireza S. Mahani, Asad Hasan, Marshall Jiang, Mansour T.A. Sharabiani

References

Mahani A.S., Hasan A., Jiang M. & Sharabiani M.T.A. (2016). Stochastic Newton Sampler: The R Package sns. Journal of Statistical Software, Code Snippets, 74(2), 1-33. doi:10.18637/jss.v074.c02

See Also

[sns.run](#)

Examples

```
## Not run:

# using RegressionFactory for generating log-likelihood and derivatives
library("RegressionFactory")

loglike.poisson <- function(beta, X, y) {
  regfac.expand.lpar(beta, X = X, y = y,
    fbase1 = fbase1.poisson.log)
}

# simulating data
K <- 5
N <- 1000
X <- matrix(runif(N * K, -0.5, +0.5), ncol = K)
beta <- runif(K, -0.5, +0.5)
y <- rpois(N, exp(X %*% beta))

beta.init <- rep(0.0, K)
beta.smp <- sns.run(beta.init, loglike.poisson,
  niter = 1000, nnr = 20, mh.diag = TRUE, X = X, y = y)

# prediction function for mean response
predmean.poisson <- function(beta, Xnew) exp(Xnew %*% beta)
ymean.new <- predict(beta.smp, predmean.poisson,
  nburnin = 100, Xnew = X)
summary(ymean.new)

# (stochastic) prediction function for response
predsmp.poisson <- function(beta, Xnew)
  rpois(nrow(Xnew), exp(Xnew %*% beta))
ysmp.new <- predict(beta.smp, predsmp.poisson,
  , nburnin = 100, Xnew = X)
summary(ysmp.new)

## End(Not run)
```

Description

SNS is a Metropolis-Hastings MCMC sampler with a multivariate Gaussian proposal function resulting from a local, second-order Taylor series expansion of log-density. The mean of the Gaussian proposal is identical to the full Newton-Raphson step from the current point. During burn-in, Newton-Raphson optimization can be performed to get close to the mode of the pdf which is unique due to convexity, resulting in faster convergence. For high dimensional densities, state space partitioning can be used to improve mixing. Support for numerical differentiation is provided using **numDeriv** package. `sns` is the low-level function for drawing one sample from the distribution. For drawing multiple samples from a (fixed) distribution, consider using `sns.run`.

Usage

```
sns(x, fghEval, rnd = TRUE, gfit = NULL, mh.diag = FALSE
    , part = NULL, numberiv = 0
    , numberiv.method = c("Richardson", "simple")
    , numberiv.args = list(), ...)
```

Arguments

<code>x</code>	Current state vector.
<code>fghEval</code>	Log-density to be sampled from. A valid log-density can have one of 3 forms: 1) return log-density, but no gradient or Hessian, 2) return a list of <code>f</code> and <code>g</code> for log-density and its gradient vector, respectively, 3) return a list of <code>f</code> , <code>g</code> , and <code>h</code> for log-density, gradient vector, and Hessian matrix. Missing derivatives are computed numerically.
<code>rnd</code>	Runs 1 iteration of Newton-Raphson optimization method (non-stochastic or 'nr' mode) when FALSE. Runs Metropolis-Hastings (stochastic or 'mcmc' mode) for drawing a sample when TRUE.
<code>gfit</code>	Gaussian fit at point <code>init</code> . If NULL then <code>sns</code> will compute a Gaussian fit at <code>x</code> .
<code>mh.diag</code>	Boolean flag, indicating whether detailed MH diagnostics such as components of acceptance test must be returned or not.
<code>part</code>	List describing partitioning of state space into subsets. Each element of the list must be an integer vector containing a set of indexes (between 1 and <code>length(x)</code> or <code>length(init)</code>) indicating which subset of all dimensions to jointly sample. These integer vectors must be mutually exclusive and collectively exhaustive, i.e. cover the entire state space and have no duplicates, in order for the partitioning to represent a valid Gibbs sampling approach. See <code>sns.make.part</code> and <code>sns.check.part</code> .
<code>numberiv</code>	Integer with value from the set $\{0, 1, 2\}$. If 0, no numerical differentiation is performed, and thus <code>fghEval</code> is expected to supply <code>f</code> , <code>g</code> and <code>h</code> . If 1, we expect <code>fghEval</code> to provide <code>f</code> and <code>g</code> , and Hessian will be calculated numerically. If

	2, <code>fghEval</code> only returns log-density, and numerical differentiation is needed to calculate gradient and Hessian.
<code>numberiv.method</code>	Method used for numeric differentiation. This is passed to the <code>grad</code> and <code>hessian</code> functions in numDeriv package. See the package documentation for details.
<code>numberiv.args</code>	Arguments to the numeric differentiation method chosen in <code>numberiv.method</code> , passed to <code>grad</code> and <code>hessian</code> functions in numDeriv . See package documentation for details.
<code>...</code>	Other arguments to be passed to <code>fghEval</code> .

Value

`sns` returns the sample drawn as a vector, with attributes:

<code>accept</code>	A boolean indicating whether the proposed point was accepted.
<code>ll</code>	Value of the log-density at the sampled point.
<code>gfit</code>	List containing Gaussian fit to pdf at the sampled point.

Note

1. Since SNS makes local Gaussian approximations to the density with the covariance matrix of the Gaussian proposal being the log-density Hessian, there is a strict requirement for the log-density to be concave.
2. Proving log-concavity for arbitrary probability distributions is non-trivial. However, distributions *generated* by replacing parameters of a concave distribution with linear expressions are known to be log-concave. This negative-definiteness invariance as well as expressions for full gradient and Hessian in terms of derivatives of low-dimensional base distributions are discussed in the vignette. The GLM expansion framework is available in the R package **RegressionFactory**.
3. See package vignette for more details on SNS theory, software, examples, and performance.

Author(s)

Alireza S. Mahani, Asad Hasan, Marshall Jiang, Mansour T.A. Sharabiani

References

- Mahani A.S., Hasan A., Jiang M. & Sharabiani M.T.A. (2016). Stochastic Newton Sampler: The R Package `sns`. *Journal of Statistical Software, Code Snippets*, 74(2), 1-33. doi:10.18637/jss.v074.c02
- Hastings, W. K. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1), 97-109.
- Qi, Y., & Minka, T. P. (2002). Hessian-based markov chain monte-carlo algorithms. 1st Cape Cod Workshop on Monte Carlo Methods.

See Also

[sns.run](#), [sns.fghEval.numaug](#)

Examples

```

## Not run:

# using RegressionFactory for generating log-likelihood and its derivatives
library(RegressionFactory)

loglike.poisson <- function(beta, X, y) {
  regfac.expand.1par(beta, X = X, y = y,
                    fbase1 = fbase1.poisson.log)
}

# simulating data
K <- 5
N <- 1000
X <- matrix(runif(N * K, -0.5, +0.5), ncol = K)
beta <- runif(K, -0.5, +0.5)
y <- rpois(N, exp(X %*% beta))

beta.init <- rep(0.0, K)

# glm estimate, for reference
beta.glm <- glm(y ~ X - 1, family = "poisson",
               start = beta.init)$coefficients

# running SNS in non-stochastic mode
# this should produce results very close to glm
beta.sns <- beta.init
for (i in 1:20)
  beta.sns <- sns(beta.sns, loglike.poisson, X = X, y = y, rnd = F)

# comparison
all.equal(as.numeric(beta.glm), as.numeric(beta.sns))

# trying numerical differentiation
loglike.poisson.fonly <- function(beta, X, y) {
  regfac.expand.1par(beta, X = X, y = y, fgh = 0,
                    fbase1 = fbase1.poisson.log)
}

beta.sns.numderiv <- beta.init
for (i in 1:20)
  beta.sns.numderiv <- sns(beta.sns.numderiv, loglike.poisson.fonly
                        , X = X, y = y, rnd = F, numderiv = 2)
all.equal(as.numeric(beta.glm), as.numeric(beta.sns.numderiv))

# add numerical derivatives to fghEval outside sns
loglike.poisson.numaug <- sns.fghEval.numaug(loglike.poisson.fonly
      , numderiv = 2)

beta.sns.numaug <- beta.init
for (i in 1:20)
  # set numderiv to 0 to avoid repeating

```



```

# numerical augmentation inside sns
beta.sns.numaug <- sns(beta.sns.numaug, loglike.poisson.numaug
                      , X = X, y = y, rnd = F, numberiv = 0)
all.equal(as.numeric(beta.glm), as.numeric(beta.sns.numaug))

## End(Not run)

```

sns.check.logdensity *Utility function for validating log-density*

Description

Utility function for validating log-density: 1) dimensional consistency of function argument, gradient and Hessian, 2) finiteness of function, gradient and Hessian, 3) closeness of analytical and numerical derivatives, and 4) negative definiteness of Hessian.

Usage

```

sns.check.logdensity(x, fghEval
                    , numberiv.method = c("Richardson", "complex")
                    , numberiv.args = list()
                    , blocks = append(list(1:length(x)), as.list(1:length(x)))
                    , dx = rep(1, length(x)), nevals = 100, negdef.tol = 1e-08, ...)
## S3 method for class 'sns.check.logdensity'
print(x, ...)

```

Arguments

x	For <code>sns.check.logdensity</code> , initial point, around which a random collection of points are generated to perform validation tests. For <code>print.sns.check.logdensity</code> , an object of class <code>sns.check.logdensity</code> , typically the output of <code>sns.check.logdensity</code> function.
fghEval	Log-density to be validated. A valid log-density can have one of 3 forms: 1) return log-density, but no gradient or Hessian, 2) return a list of f and g for log-density and its gradient vector, respectively, 3) return a list of f, g, and h for log-density, gradient vector, and Hessian matrix.
numberiv.method	Method used for numeric differentiation. This is passed to the <code>grad</code> and <code>hessian</code> functions in numDeriv package. See the package documentation for details.
numberiv.args	Arguments to the numeric differentiation method chosen in <code>numberiv.method</code> , passed to <code>grad</code> and <code>hessian</code> functions in numDeriv . See package documentation for details.
blocks	A list of state space subsets (identified by their positional indexes), for which negative-definiteness of Hessian blocks are to be tested. The default is to test for 1) entire state space, and 2) each dimension individually.

dx	A vector of same length as x . For i 'th dimension, <code>nevals</code> values are sampled from a uniform distribution with min/max values equal to $x[i]-0.5*dx[i]$ and $x[i]+0.5*dx[i]$, respectively. Vectors smaller than <code>length(x)</code> are extended as needed by recycling the provided values for <code>dx</code> .
nevals	Number of points in state space, for which validation tests will be performed.
negdef.tol	Lower bound for absolute value of (negative) eigenvalues of Hessian, evaluated at each of the <code>nevals</code> points in the state space. If one or more eigenvalues have absolute values smaller than <code>negdef.tol</code> , log-density is declared non-log-concave at that point.
...	Other arguments to be passed to <code>fghEval</code> .

Value

`sns.check.logdensity` returns a list of class `sns.check.logdensity`, with the following elements:

<code>check.ld.struct</code>	Boolean flag, indicating whether log-density <code>fghEval</code> has one of the 3 forms of output, described above.
<code>numberiv</code>	Integer with values of 0, 1, 2. A value of 0 means analytical gradient and Hessian have been provided, and thus there is no need for numerical differentiation. 1 means analytical gradient is provided, but Hessian must be calculated numerically. 2 means both gradient and Hessian must be numerically calculated. Users can pass this value to subsequent <code>sns</code> or <code>sns.run</code> calls.
<code>check.length.g</code>	Boolean flag, indicating whether length of gradient vector (element <code>g</code>) returned by <code>fghEval</code> equals <code>length(x)</code> .
<code>check.dim.h</code>	Boolean flag, indicating whether number of rows and columns of the Hessian matrix (element <code>h</code>) returned by <code>fghEval</code> equal <code>length(x)</code> .
<code>x.mat</code>	Collection of state space vectors (one per row), for which validation tests are performed. It has <code>nevals</code> rows and <code>length(x)</code> columns.
<code>t.ivals</code>	Time spent on evaluating <code>fghEval</code> on <code>nevals</code> points chosen randomly in the neighborhood of x , as specified by <code>dx</code> . This includes log-density and, if provided, analytical evaluations of gradient and Hessian.
<code>t.num.ivals</code>	Time spent on evaluating the numeric version of <code>fghEval</code> , in which gradient and Hessian are computed numerically, using <code>grad</code> and <code>hessian</code> functions in the numDeriv package. Comparison of this number with <code>t.ivals</code> provides the user with insight into the relative speed of numerical differentiation compared to analytical versions.
<code>f.vec</code>	Vector of log-density values for state space vectors listed in <code>x.mat</code> .
<code>g.mat.num</code>	Collection of numerically-computed gradient vectors for state space values listed in <code>x.mat</code> , with the same dimension conventions.
<code>is.g.num.finite</code>	Boolean flag, indicating whether all numerically-computed gradient vectors have finite values.

h.array.num	Collection of numerically-computed Hessian matrices at points listed in <code>x.mat</code> . First dimension is of length <code>nevals</code> , and the remaining two dimensions equal <code>length(x)</code> .
is.h.num.finite	Boolean flag, indicating whether all numerically-computed Hessian matrices have finite values.
g.mat	Collection of analytically-computed gradient vectors for state space values listed in <code>x.mat</code> , with the same dimension conventions. This is only available if <code>fghEval</code> has a <code>g</code> field; otherwise NA.
is.g.finite	Boolean flag (if available), indicating whether all analytically-computed gradient vectors have finite values (if available).
g.diff.max	If available, maximum relative difference between analytical and numerical gradient vectors, over all <code>nevals</code> points in <code>x.mat</code> . Relative difference is defined as L2 norm of difference between the two gradient vectors, divided by the L2 norm of the analytical gradient vector.
h.array	If available, collection of analytically-computed Hessian matrices at points listed in <code>x.mat</code> . Dimensional conventions are the same as <code>h.array.num</code> .
is.h.finite	Boolean flag (if available), indicating whether all analytically-computed Hessian matrices have finite values.
h.diff.max	If available, maximum relative difference between analytical and numerical Hessian matrices, over all <code>nevals</code> points in <code>x.mat</code> . Relative difference is defined as the Frobenius norm of difference of analytical and numerical Hessian matrices, divided by the Frobenius norm of analytical Hessian.
is.negdef.num	Boolean flag, indicating whether numerical Hessian is negative-definite at all state space points indicated in <code>x.mat</code> .
is.negdef	Boolean flag, indicating whether analytical Hessian is negative-definite at all state space points indicated in <code>x.mat</code> .

Note

1. Validation tests performed in `sns.check.logdensity` cannot prove that a log-density is twice-differentiable, or globally concave. However, when e.g. log-density Hessian is seen to be non-negative-definite at one of the points tested, we can definitively say that the Hessian is not globally negative-definite, and therefore `sns` should not be used for sampling from this distribution. Users must generally consider this function as a supplement to analytical work.
2. See package vignette for more details on SNS theory, software, examples, and performance.

Author(s)

Alireza S. Mahani, Asad Hasan, Marshall Jiang, Mansour T.A. Sharabiani

References

Mahani A.S., Hasan A., Jiang M. & Sharabiani M.T.A. (2016). Stochastic Newton Sampler: The R Package `sns`. *Journal of Statistical Software, Code Snippets*, 74(2), 1-33. doi:10.18637/jss.v074.c02

Examples

```

## Not run:

# using RegressionFactory for generating log-likelihood and its derivatives
library(RegressionFactory)

loglike.poisson <- function(beta, X, y) {
  regfac.expand.1par(beta, X = X, y = y,
                    fbase1 = fbase1.poisson.log)
}

# simulating data
K <- 5
N <- 1000
X <- matrix(runif(N * K, -0.5, +0.5), ncol = K)
beta <- runif(K, -0.5, +0.5)
y <- rpois(N, exp(X %*% beta))

beta.init <- rep(0.0, K)

my.check <- sns.check.logdensity(beta.init, loglike.poisson
, X = X, y = y, blocks = list(1:K))
my.check

# mistake in log-likelihood gradient
loglike.poisson.wrong <- function(beta, X, y) {
  ret <- regfac.expand.1par(beta, X = X, y = y,
                          fbase1 = fbase1.poisson.log)

  ret$g <- 1.2 * ret$g
  return (ret)
}
# maximum relative diff in gradient is now much larger
my.check.wrong <- sns.check.logdensity(beta.init
, loglike.poisson.wrong, X = X, y = y, blocks = list(1:K))
my.check.wrong

# mistake in log-likelihood Hessian
loglike.poisson.wrong.2 <- function(beta, X, y) {
  ret <- regfac.expand.1par(beta, X = X, y = y,
                          fbase1 = fbase1.poisson.log)

  ret$h <- 1.2 * ret$h
  return (ret)
}
# maximum relative diff in Hessian is now much larger
my.check.wrong.2 <- sns.check.logdensity(beta.init
, loglike.poisson.wrong.2, X = X, y = y, blocks = list(1:K))
my.check.wrong.2

## End(Not run)

```

sns.fghEval.numaug	<i>Utility function for augmentation of a log-density function with numerical gradient and Hessian as needed</i>
--------------------	--

Description

Augmenting a log-density with numerical gradient and Hessian, so it can be used by `sns` or `sns.run`. This augmentation will also be done inside the function, if the value of `numderiv` parameter passed to `sns` and `sns.run` is 1 or 2. The advantage of using `sns.fghEval.numaug` outside these functions is efficiency, since the augmentation code will not have to be executed in every function call. Users must set `numderiv` to 0 when calling `sns` or `sns.run` if calling `sns.fghEval.numaug` first. See example.

Usage

```
sns.fghEval.numaug(fghEval, numderiv = 0
, numderiv.method = c("Richardson", "simple")
, numderiv.args = list())
```

Arguments

<code>fghEval</code>	Log-density to be sampled from. A valid log-density can have one of 3 forms: 1) return log-density, but no gradient or Hessian, 2) return a list of <code>f</code> and <code>g</code> for log-density and its gradient vector, respectively, 3) return a list of <code>f</code> , <code>g</code> , and <code>h</code> for log-density, gradient vector, and Hessian matrix. Missing derivatives are computed numerically.
<code>numderiv</code>	This must be matched with <code>fghEval</code> : Integer with value from the set 0, 1, 2. If 0, no numerical differentiation is performed, and thus <code>fghEval</code> is expected to supply <code>f</code> , <code>g</code> and <code>h</code> . If 1, we expect <code>fghEval</code> to provide <code>f</code> and <code>g</code> , and Hessian will be calculated numerically. If 2, <code>fghEval</code> only returns log-density, and numerical differentiation is needed to calculate gradient and Hessian.
<code>numderiv.method</code>	Method used for numeric differentiation. This is passed to the <code>grad</code> and <code>hessian</code> functions in numDeriv package. See the package documentation for details.
<code>numderiv.args</code>	Arguments to the numeric differentiation method chosen in <code>numderiv.method</code> , passed to <code>grad</code> and <code>hessian</code> functions in numDeriv . See package documentation for details.

Value

A function, accepting same arguments as `fghEval`, but guaranteed to return the original log-density, plus gradient and Hessian (both of which could possibly be numerically calculated). If `numderiv=0`, `fghEval` is returned without change. The function will return log-density, gradient and Hessian as elements `f`, `g` and `h` of a list.

Note

See package vignette for more details on SNS theory, software, examples, and performance.

Author(s)

Alireza S. Mahani, Asad Hasan, Marshall Jiang, Mansour T.A. Sharabiani

References

Mahani A.S., Hasan A., Jiang M. & Sharabiani M.T.A. (2016). Stochastic Newton Sampler: The R Package sns. *Journal of Statistical Software, Code Snippets*, 74(2), 1-33. doi:10.18637/jss.v074.c02

See Also

[sns](#), [sns.run](#)

sns.make.part

Utility Functions for Creating and Validating State Space Partitions

Description

Utility functions for creating and validating state space partitions, to be used in SNS for improving the mixing of sampled chains for high-dimensional posteriors.

Usage

```
sns.make.part(K, nsubset, method = "naive")
sns.check.part(part, K)
```

Arguments

K	Dimensionality of state space.
nsubset	Number of subsets to partition the state space dimensions into.
method	Method used for state space partitioning. Currently, only naive method is implemented, where coordinates are distributed evenly (or as evenly as possible) across subsets.
part	A list of length nsubset, with each element a vector of integer values, representing the coordinates belonging to a subset. This list is the output of sns.make.part.

Value

sns.make.part produces a list of integer vectors, each containing coordinates belonging to the same subset. sns.check.part produces a boolean flag, indicating whether or not the partition list is valid or not. The subset members must constitute a mutually-exclusive, collectively-exhaustive set relative to 1:K.

Author(s)

Alireza S. Mahani, Asad Hasan, Marshall Jiang, Mansour T.A. Sharabiani

References

Mahani A.S., Hasan A., Jiang M. & Sharabiani M.T.A. (2016). Stochastic Newton Sampler: The R Package sns. Journal of Statistical Software, Code Snippets, 74(2), 1-33. doi:10.18637/jss.v074.c02

See Also

[sns](#), [sns.run](#)

Examples

```
# creating a valid partition of a 6-dimensional state space
my.part.valid <- list(c(1,2,3), c(4,5,6))
is.valid.1 <- sns.check.part(my.part.valid, 6)
cat("is partition valid: ", is.valid.1, "\n")

# creating an invalid partition of a 6-dimensional state space
# (coordinate 4 is missing)
my.part.invalid <- list(c(1,2,3), c(5,6))
is.valid.2 <- sns.check.part(my.part.invalid, 6)
cat("is partition valid: ", is.valid.2, "\n")
```

sns.run

Drawing multiple samples using Stochastic Newton Sampler

Description

This is a wrapper around `sns`, allowing one to draw multiple samples from a distribution while collecting diagnostic information.

Usage

```
sns.run(init, fghEval, niter = 100, nnr = min(10, round(niter/4))
, mh.diag = FALSE, part = NULL, print.level = 0
, report.progress = ceiling(niter/10)
, numderiv = 0, numderiv.method = c("Richardson", "simple")
, numderiv.args = list()
, ...)
```

Arguments

<code>init</code>	Initial value for the MCMC chain.
<code>fghEval</code>	Log-density to be sampled from. A valid log-density can have one of 3 forms: 1) return log-density, but no gradient or Hessian, 2) return a list of f and g for log-density and its gradient vector, respectively, 3) return a list of f, g, and h for log-density, gradient vector, and Hessian matrix. Missing derivatives are computed numerically.
<code>niter</code>	Number of iterations to perform (in ‘nr’ and ‘mcmc’ mode combined).
<code>nnr</code>	Number of initial iterations to spend in ‘nr’ mode.
<code>mh.diag</code>	Boolean flag, indicating whether detailed MH diagnostics such as components of acceptance test must be returned or not.
<code>part</code>	List describing partitioning of state space into subsets. Each element of the list must be an integer vector containing a set of indexes (between 1 and <code>length(x)</code> or <code>length(init)</code>) indicating which subset of all dimensions to jointly sample. These integer vectors must be mutually exclusive and collectively exhaustive, i.e. cover the entire state space and have no duplicates, in order for the partitioning to represent a valid Gibbs sampling approach. See <code>sns.make.part</code> and <code>sns.check.part</code> .
<code>print.level</code>	If greater than 0, print sampling progress report.
<code>report.progress</code>	Number of sampling iterations to wait before printing progress reports.
<code>numberiv</code>	Integer with value from the set $\{0, 1, 2\}$. If 0, no numerical differentiation is performed, and thus <code>fghEval</code> is expected to supply f, g and h. If 1, we expect <code>fghEval</code> to provide f and g, and Hessian will be calculated numerically. If 2, <code>fghEval</code> only returns log-density, and numerical differentiation is needed to calculate gradient and Hessian.
<code>numberiv.method</code>	Method used for numeric differentiation. This is passed to the <code>grad</code> and <code>hessian</code> functions in numDeriv package. See the package documentation for details.
<code>numberiv.args</code>	Arguments to the numeric differentiation method chosen in <code>numberiv.method</code> , passed to <code>grad</code> and <code>hessian</code> functions in numDeriv . See package documentation for details.
<code>...</code>	Other parameters to be passed to <code>fghEval</code> .

Value

`sns.run` returns an object of class `sns` with elements:

<code>samplesMat</code>	A matrix object with <code>nsample</code> rows and <code>K</code> cols.
<code>acceptance</code>	Metropolis proposal percentage acceptance.
<code>burn.iters</code>	Number of burn-in iterations.
<code>sample.time</code>	Time in seconds spent in sampling.
<code>burnin.time</code>	Time in seconds spent in burn-in.

Note

1. `sns.run` cannot be used if SNS is being run as part of a Gibbs cycle, such that the conditional distribution being sampled by SNS changes from one iteration to next. In such cases, `sns` must be used instead, inside an explicit Gibbs-cycle for loop.
2. See package vignette for more details on SNS theory, software, examples, and performance.

Author(s)

Alireza S. Mahani, Asad Hasan, Marshall Jiang, Mansour T.A. Sharabiani

References

Mahani A.S., Hasan A., Jiang M. & Sharabiani M.T.A. (2016). Stochastic Newton Sampler: The R Package `sns`. *Journal of Statistical Software, Code Snippets*, 74(2), 1-33. doi:10.18637/jss.v074.c02

See Also

`sns`, [summary.sns](#), [plot.sns](#), [predict.sns](#)

Examples

```
## Not run:

# using RegressionFactory for generating log-likelihood and its derivatives
library(RegressionFactory)

loglike.poisson <- function(beta, X, y) {
  regfac.expand.lpar(beta, X = X, y = y,
    fbase1 = fbase1.poisson.log)
}

# simulating data
K <- 5
N <- 1000
X <- matrix(runif(N * K, -0.5, +0.5), ncol = K)
beta <- runif(K, -0.5, +0.5)
y <- rpois(N, exp(X))

beta.init <- rep(0.0, K)

# glm estimate (ML), for reference
beta.glm <- glm(y ~ X - 1, family = "poisson",
  start = beta.init)$coefficients

# sampling of likelihood
beta.smp <- sns.run(init = beta.init
  , fghEval = loglike.poisson, niter = 1000
  , nnr = 20, X = X, y = y)
smp.summ <- summary(beta.smp)

# compare mean of samples against ML estimate (from glm)
```

```

cbind(beta.glm, smp.summ$smp$mean)

# trying numerical differentiation
loglike.poisson.fonly <- function(beta, X, y) {
  regfac.expand.lpar(beta, X = X, y = y, fgh = 0,
    fbase1 = fbase1.poisson.log)
}
beta.smp <- sns.run(init = beta.init
  , fghEval = loglike.poisson.fonly, niter = 1000, nnr = 20
  , X = X, y = y, numberiv = 2)
smp.summ <- summary(beta.smp)
cbind(beta.glm, smp.summ$smp$mean)

## End(Not run)

```

summary.sns

Summarizing "sns" Objects

Description

Methods for summarizing the output of `sns.run`, and for printing the summary.

Usage

```

## S3 method for class 'sns'
summary(object, quantiles = c(0.025, 0.5, 0.975)
  , pval.ref = 0.0, nburnin = max(nrow(object)/2, attr(object, "nnr"))
  , end = nrow(object), thin = 1, ess.method = c("coda", "ise"), ...)
## S3 method for class 'summary.sns'
print(x, ...)

```

Arguments

<code>object</code>	An object of class "sns", typically the output of <code>sns.run</code> .
<code>quantiles</code>	Values for which sample-based quantiles are calculated.
<code>pval.ref</code>	Reference value for state space variables, used for calculating sample-based p-values.
<code>nburnin</code>	Number of initial iterations to discard before calculating the sample statistics. A warning is issued if this number is smaller than the initial iterations run in NR mode.
<code>end</code>	Last iteration to use for calculating sample statistics. Defaults to last iteration.
<code>thin</code>	One out of thin samples are kept for calculating sample statistics. Default is 1, using all samples within specified range.
<code>ess.method</code>	Method used for calculating effective sample size. Default is to call <code>effectiveSize</code> from package <code>coda</code> .
<code>x</code>	An object of class "summary.sns", typically the output of <code>summary.sns</code> .
<code>...</code>	Arguments passed to/from other functions.

Value

summary.sns returns a list with these elements:

K	Dimensionality of state space.
nnr	Number of NR (Newton-Raphson) iterations performed at the beginning.
nburnin	Number of burn-in iterations. These are discarded before calculating sample statistics.
end	Last iteration to use for calculating sample statistics.
thin	One out of every thin iterations within the specified range is used for calculating sample statistics.
niter	Total iterations, including NR and MCMC modes.
nsmpl	Number of samples within specified range (before applying thinning).
nseq	Number of samples used for calculating sample statistics (after applying thinning).
npart	Number of subsets used in state space partitioning. If no partitioning is done, the value is 1.
accept.rate	Acceptance rate for the MH transition proposals, calculated over nsmpl iterations.
reldev.mean	Mean relative deviation from quadratic approximation, defined as difference between actual log-density change and the value predicted from quadratic fit at density maximum, divided by the actual change. The location of density maximum is assumed to be the value at the end of the last NR iteration. Therefore, for this measure to be accurate, users must ensure nnr is sufficiently large to allow for convergence of the optimization phase.
pval.ref	Same as input.
ess.method	Same as input.
smp	A list with elements mean, sd, ess, quantiles, pval representing sample-based mean, standard deviation, effective size, quantiles and sample-based p-values, based on specified range and using thinning (if specified).

Author(s)

Alireza S. Mahani, Asad Hasan, Marshall Jiang, Mansour T.A. Sharabiani

References

Mahani A.S., Hasan A., Jiang M. & Sharabiani M.T.A. (2016). Stochastic Newton Sampler: The R Package sns. *Journal of Statistical Software, Code Snippets*, 74(2), 1-33. doi:10.18637/jss.v074.c02

See Also

[sns.run](#)

Index

* **Metropolis**

sns, 6
sns.run, 15

* **mcmc**

sns, 6
sns.run, 15

* **multivariate**

sns, 6
sns.run, 15

* **sampling**

sns, 6
sns.run, 15

ess, 2

plot.sns, 3, 17

predict.sns, 4, 17

print.sns.check.logdensity
(sns.check.logdensity), 9

print.summary.predict.sns
(predict.sns), 4

print.summary.sns (summary.sns), 18

sns, 6, 14, 15

sns.check.logdensity, 9

sns.check.part (sns.make.part), 14

sns.fghEval.numaug, 7, 13

sns.make.part, 14

sns.run, 3–5, 7, 14, 15, 15, 18, 19

summary.predict.sns (predict.sns), 4

summary.sns, 17, 18