

Package ‘dapper’

September 6, 2026

Title Data Augmentation for Private Posterior Estimation

Version 1.1.1

Description A data augmentation based sampler for conducting privacy-aware Bayesian inference. The `dapper_sample()` function takes an existing sampler as input and automatically constructs a privacy-aware sampler. The process of constructing a sampler is simplified through the specification of four independent modules, allowing for easy comparison between different privacy mechanisms by only swapping out the relevant modules. Probability mass functions for the discrete Gaussian and discrete Laplacian are provided to facilitate analyses dealing with privatized count data. The output of `dapper_sample()` can be analyzed using many of the same tools from the 'rstan' ecosystem. For methodological details on the sampler see Ju et al. (2022) <[doi:10.48550/arXiv.2206.00710](https://doi.org/10.48550/arXiv.2206.00710)>, and for details on the discrete Gaussian and discrete Laplacian distributions see Canonne et al. (2020) <[doi:10.48550/arXiv.2004.00010](https://doi.org/10.48550/arXiv.2004.00010)>.

License MIT + file LICENSE

Encoding UTF-8

RoxygenNote 7.3.3

URL <https://github.com/mango-empire/dapper>

BugReports <https://github.com/mango-empire/dapper/issues>

Suggests ggplot2, knitr, rmarkdown, testthat (>= 3.0.0)

VignetteBuilder knitr

Config/testthat/edition 3

Imports bayesplot, checkmate, furrr, memoise, posterior, progressr, stats

NeedsCompilation no

Author Kevin Eng [aut, cre, cph]

Maintainer Kevin Eng <kevine1221@gmail.com>

Repository CRAN

Date/Publication 2026-09-06 13:00:10 UTC

Contents

dapper_sample	2
ddlaplace	4
ddnorm	5
new_privacy	6
plot.dpout	8
summary.dpout	8

Index	9
--------------	----------

dapper_sample	<i>Private Posterior Sampler</i>
---------------	----------------------------------

Description

Generates samples from the private posterior using a data augmentation framework.

Usage

```
dapper_sample(
  data_model = NULL,
  sdp = NULL,
  init_par = NULL,
  seed = NULL,
  niter = 2000,
  warmup = floor(niter/2),
  chains = 1
)
```

Arguments

data_model	a data model represented by a privacy class object.
sdp	the observed privatized data. Must be a vector or matrix.
init_par	initial starting point of the chain.
seed	set random seed.
niter	number of draws.
warmup	number of iterations to discard as warmup. Default is half of niter.
chains	number of MCMC chains to run. Can be done in parallel or sequentially.

Details

Generates samples from the private posterior implied by data_model. The data_model input must be an object of class privacy which is created using the new_privacy() constructor. MCMC chains can be run in parallel using furrr::future_map(). See the **furrr** package documentation for specifics. Long computations can be monitored with the **progressr** package.

Value

A dpout object which contains:

- chain: a `draws_matrix` object containing chains * (niter - warmup) draws from the private posterior.
- mean_accept: a (niter - warmup) row matrix containing the average acceptance rate over all latent records for each iteration. Each column corresponds to a chain.
- comp_accept: An n-by-chains matrix, where n is the number of latent records. Each entry is the mean acceptance rate for one record in one chain over the retained post-warmup iterations.

References

Ju, N., Awan, J. A., Gong, R., & Rao, V. A. (2022). Data Augmentation MCMC for Bayesian Inference from Privatized Data. *arXiv*. doi:[10.48550/ARXIV.2206.00710](https://doi.org/10.48550/ARXIV.2206.00710)

See Also

`new_privacy()`

Examples

```
#simulate confidential data
# The privacy mechanism adds Gaussian noise to the sample mean.
set.seed(1)
n <- 100
eps <- 3
y <- rnorm(n, mean = -2, sd = 1)
sdp <- mean(y) + rnorm(1, 0, 1/eps)

posterior_f <- function(dmat, theta) {
  x <- c(dmat)
  xbar <- mean(x)
  n <- length(x)
  pr_m <- 0
  pr_s2 <- 4
  ps_s2 <- 1/(1/pr_s2 + n)
  ps_m <- ps_s2 * ((1/pr_s2)*pr_m + n * xbar)
  rnorm(1, mean = ps_m, sd = sqrt(ps_s2))
}
latent_f <- function(theta) {
  matrix(rnorm(100, mean = theta, sd = 1), ncol = 1)
}
statistic_f <- function(xi, sdp, i) {
  xi
}
mechanism_f <- function(sdp, sx) {
  sum(dnorm(sdp - sx/n, 0, 1/eps, TRUE))
}
dmod <- new_privacy(posterior_f = posterior_f,
  latent_f = latent_f,
  mechanism_f = mechanism_f,
```

```

    statistic_f = statistic_f,
    npar = 1)

out <- dapper_sample(dmod,
                    sdp = sdp,
                    init_par = -2,
                    niter = 500)

summary(out)
plot(out)

# for parallel computing we 'plan' a session
# the code below uses 2 CPU cores for parallel computing
library(furrr)
plan(multisession, workers = 2)
out <- dapper_sample(dmod,
                    sdp = sdp,
                    init_par = -2,
                    niter = 500,
                    chains = 2)

# to go back to sequential computing we use
plan(sequential)

```

ddlaplace

Discrete Laplace Distribution

Description

The probability mass function and random number generator for the discrete Laplacian distribution.

Usage

```
ddlaplace(x, scale = 1, log = FALSE)
```

```
rdlaplace(n, scale = 1)
```

Arguments

x	a vector of quantiles.
scale	the scale parameter. A positive number for ddlaplace() and a positive integer for rdlaplace().
log	logical; if TRUE, probabilities are given as log(p).
n	number of random deviates.

Details

Probability mass function

$$P[X = x] = \frac{e^{1/t} - 1}{e^{1/t} + 1} e^{-|x|/t}.$$

Value

- `ddlaplace()` returns a numeric vector representing the probability mass function of the discrete Laplace distribution.
- `rdlaplace()` returns a numeric vector of random samples from the discrete Laplace distribution.

References

Canonne, C. L., Kamath, G., & Steinke, T. (2020). The Discrete Gaussian for Differential Privacy. *arXiv*. doi:10.48550/ARXIV.2004.00010

Examples

```
# mass function
ddlaplace(0)

# mass function is vectorized
ddlaplace(0:10, scale = 5)

# generate random samples
rdlaplace(10)
```

ddnorm

The Discrete Gaussian Distribution

Description

The probability mass function and random number generator for the discrete Gaussian distribution with mean μ and scale parameter σ .

Usage

```
ddnorm(x, mu = 0, sigma = 1, log = FALSE)

rdnorm(n, mu = 0, sigma = 1)
```

Arguments

<code>x</code>	vector of quantiles.
<code>mu</code>	location parameter. Must be an integer for both <code>ddnorm()</code> and <code>rdnorm()</code> .
<code>sigma</code>	scale parameter.
<code>log</code>	logical; if TRUE, log unnormalized probabilities are returned.
<code>n</code>	number of random deviates.

Details

Probability mass function

$$P[X = x] = \frac{e^{-(x-\mu)^2/2\sigma^2}}{\sum_{y \in \mathbb{Z}} e^{-(y-\mu)^2/2\sigma^2}}.$$

Value

- `ddnorm()` returns a numeric vector representing the probability mass function of the discrete Gaussian distribution.
- `rdnorm()` returns a numeric vector of random samples from the discrete Gaussian distribution.

References

Canonne, C. L., Kamath, G., & Steinke, T. (2020). The Discrete Gaussian for Differential Privacy. *arXiv*. doi:[10.48550/ARXIV.2004.00010](https://doi.org/10.48550/ARXIV.2004.00010)

Examples

```
# mass function
ddnorm(0)

# mass function is also vectorized
ddnorm(0:10, mu = 0, sigma = 5)

# generate random samples
rdnorm(10)
```

new_privacy

privacy *Object Constructor*.

Description

Creates a privacy object to be used as input into `dapper_sample()`.

Usage

```
new_privacy(
  posterior_f = NULL,
  latent_f = NULL,
  mechanism_f = NULL,
  statistic_f = NULL,
  npar = NULL,
  varnames = NULL
)
```

Arguments

<code>posterior_f</code>	a function that draws posterior samples given the confidential data.
<code>latent_f</code>	A function with signature <code>latent_f(theta)</code> that generates a latent confidential dataset conditional on <code>theta</code> .
<code>mechanism_f</code>	a function that represents the log likelihood of the privacy mechanism.
<code>statistic_f</code>	a function that calculates the statistic to be released.
<code>npar</code>	dimension of the parameter being estimated.
<code>varnames</code>	an optional character vector of parameter names. Used to label summary outputs.

Details

- `posterior_f()` is a function which makes draws from the posterior sampler. It has the syntax `posterior_f(dmat, theta)`. Here `dmat` is a numeric matrix representing the confidential database and `theta` is a numeric vector which serves as the initialization point for a one sample draw. The easiest, bug-free way to construct `posterior_f()` is to use a conjugate prior. However, this function can also be constructed by wrapping a MCMC sampler generated from other R packages (e.g. **rstan**, **fmcme**, **adaptMCMC**).
- `latent_f(theta)` must take exactly one argument named `theta`, a numeric vector of length `npar`. It must return a matrix with one confidential record per row and one data variable per column. The records (rows) must be independent and identically distributed conditional on `theta`. Variables within a record may be dependent. The returned matrix must have the same dimensions on every call. For univariate data, return a one-column matrix.
- `mechanism_f()` is a function that represents the log of the privacy mechanism density. This function has the form `mechanism_f(sdp, sx)` where `sdp` and `sx` are both either a numeric vector or matrix. The arguments must appear in the exact stated order with the same variables names as mentioned. Finally, the return value of `mechanism_f()` must be a numeric vector of length one.
- `statistic_f(xi, sdp, i)` computes one record's contribution to the statistic used by the privacy mechanism. `xi` is row `i` of the latent confidential dataset, supplied as a vector. `sdp` is the observed privatized release, and `i` is the row index. Argument names and order must match this signature.
The sampler sums these contributions over all records and passes the result as `sx` to `mechanism_f(sdp, sx)`. Return a numeric vector when `sdp` is a numeric vector, or a matrix when `sdp` is a matrix. Contributions must have consistent lengths or dimensions across records.
- `npar` is an integer equal to the dimension of `theta`.

Value

A S3 object of class `privacy`.

plot.dpout	<i>Plot dpout object.</i>
------------	---------------------------

Description

Plot dpout object.

Usage

```
## S3 method for class 'dpout'  
plot(x, ...)
```

Arguments

x	dp_out object.
...	optional arguments to mcmc_trace().

Value

trace plots.

summary.dpout	<i>Summarise dpout object.</i>
---------------	--------------------------------

Description

Summarise dpout object.

Usage

```
## S3 method for class 'dpout'  
summary(object, ...)
```

Arguments

object	dp_out object
...	optional arguments to summarise_draws().

Value

A [draws_summary](#) tibble containing summary statistics (such as mean, sd, ess_bulk, and rhat).
See [posterior::summarise_draws\(\)](#) for details on how these statistics are calculated.

Index

dapper_sample, [2](#)
ddlaplace, [4](#)
ddnorm, [5](#)
draws_matrix, [3](#)
draws_summary, [8](#)

new_privacy, [6](#)
new_privacy(), [3](#)

plot.dpout, [8](#)
posterior::summarise_draws(), [8](#)

rdlaplace (ddlaplace), [4](#)
rddnorm (ddnorm), [5](#)

summary.dpout, [8](#)