

Package ‘deepgp’

October 29, 2020

Type Package

Title Sequential Design for Deep Gaussian Processes using MCMC

Version 0.1.0

Date 2020-10-8

Author Annie Sauer <anniees@vt.edu>

Maintainer Annie Sauer <anniees@vt.edu>

Depends R (>= 3.6)

Description Performs model fitting and sequential design for deep Gaussian processes using MCMC and elliptical slice sampling. Models extend up to three layers deep; a one layer model is equivalent to typical Gaussian process regression. Sequential design criteria include integrated mean square prediction error (IMSPE), active learning Cohn (ALC), and expected improvement (EI). Covariance structure is based on inverse exponentiated squared euclidean distance. Applicable to noisy and deterministic functions. Incorporates SNOW parallelization and utilizes C under the hood. Manuscript forthcoming; see Damianou and Lawrence (2013) <arXiv:1211.0358> for deep Gaussian process models and Murray, Adams, and MacKay (2010) <arXiv:1001.0175> for elliptical slice sampling.

License LGPL

Encoding UTF-8

NeedsCompilation yes

Imports grDevices, graphics, stats, doParallel, foreach, parallel

Suggests akima, knitr

RoxygenNote 7.1.1

Repository CRAN

Date/Publication 2020-10-29 10:20:08 UTC

R topics documented:

deepgp-package	2
ALC	5

calc_K	7
continue	8
EI	9
fit_one_layer	10
fit_three_layer	12
fit_two_layer	15
IMSPE	18
plot	20
predict	21
rmse	23
score	23
sq_dist	24
trim	25

Index	26
--------------	-----------

deepgp-package	<i>Package deepgp</i>
----------------	-----------------------

Description

Performs model fitting and sequential design for deep Gaussian processes using MCMC and elliptical slice sampling. Models extend up to three layers deep; a one layer model is equivalent to typical Gaussian process regression. Sequential design criteria include integrated mean square prediction error (IMSPE), active learning Cohn (ALC), and expected improvement (EI). Covariance structure is based on inverse exponentiated squared euclidean distance. Applicable to noisy and deterministic functions. Incorporates SNOW parallelization and utilizes C under the hood. Manuscript forthcoming; see Damianou and Lawrence (2013) <arXiv:1211.0358> for deep Gaussian process models and Murray, Adams, and MacKay (2010) <arXiv:1001.0175> for elliptical slice sampling.

Important Functions

- [fit_one_layer](#): conducts MCMC sampling of hyperparameters for a one layer GP
- [fit_two_layer](#): conducts MCMC sampling of hyperparameters for a two layer deep GP
- [fit_three_layer](#): conducts MCMC sampling of hyperparameters for a three layer deep GP
- [continue](#): collects additional MCMC samples
- [trim](#): cuts off burn-in and optionally thins samples
- [predict](#): calculates posterior mean and covariance over a set of input locations
- [plot](#): produces trace plots, hidden layer plots, and posterior plots
- [ALC](#): calculates active learning Cohn over set of input locations using reference grid
- [IMSPE](#): calculates integrated mean square prediction error over set of input locations
- [EI](#): calculates expected improvement over set of input locations

Author(s)

Annie Sauer <anniees@vt.edu>

References

- Binois, M, J Huang, RB Gramacy, and M Ludkovski. 2019. "Replication or Exploration? Sequential Design for Stochastic Simulation Experiments." *Technometrics* 61, 7-23. Taylor & Francis. doi:10.1080/00401706.2018.1469433.
- Damianou, A and N Lawrence. (2013). "Deep gaussian processes." *Artificial Intelligence and Statistics*, 207-215.
- Gramacy, RB. *Surrogates: Gaussian Process Modeling, Design, and Optimization for the Applied Sciences*. Chapman Hall, 2020.
- Jones, DR, M Schonlau, and WJ Welch. 1998. "Efficient Global Optimization of Expensive Black-Box Functions." *Journal of Global Optimization* 13, 455-492. doi:10.1023/A:1008306431147.
- Murray, I, RP Adams, and D MacKay. 2010. "Elliptical slice sampling." *Journal of Machine Learning Research* 9, 541-548.
- Seo, S, M Wallat, T Graepel, and K Obermayer. 2000. "Gaussian Process Regression: Active Data Selection and Test Point Rejection." In *Mustererkennung 2000*, 27–34. New York, NY: Springer-Verlag.

Examples

```
# 1. One Layer and EI -----
f <- function(x) {
  sin(5 * pi * x) / (2 * x) + (x - 1) ^ 4
}

# Training data
x <- seq(0.5, 2, length = 30)
y <- f(x) + rnorm(30, 0, 0.01)

# Testing data
xx <- seq(0.5, 2, length = 100)
yy <- f(xx)

# Standardize inputs and outputs
xx <- (xx - min(x)) / (max(x) - min(x))
x <- (x - min(x)) / (max(x) - min(x))
yy <- (yy - mean(y)) / sd(y)
y <- (y - mean(y)) / sd(y)

# Conduct MCMC
fit <- fit_one_layer(x, y, ncmc = 10000)
plot(fit) # investigate trace plots
fit <- trim(fit, 8000, 2)

# Predict and calculate EI
```

```

fit <- predict(fit, xx, lite = FALSE)
ei <- EI(fit)

# Visualize Fit
plot(fit)
par(new = TRUE) # overlay EI
plot(xx, ei$value, type = 'l', lty = 2, axes = FALSE, xlab = '', ylab = '')

# Select next design point
x_new <- xx[which.max(ei$value)]

# Evaluate fit
rmse(yy, fit$mean) # lower is better
score(yy, fit$mean, fit$Sigma) # higher is better

# 2. Two Layer and ALC -----

f <- function(x) {
  exp(-10 * x) * (cos(10 * pi * x - 1) + sin(10 * pi * x - 1)) * 5 - 0.2
}

# Training data
x <- seq(0, 1, length = 30)
y <- f(x) + rnorm(30, 0, 0.05)

# Testing data
xx <- seq(0, 1, length = 100)
yy <- f(xx)

# Conduct MCMC
fit <- fit_two_layer(x, y, D = 1, nmcmc = 9000)
fit <- continue(fit, 1000)
plot(fit) # investigate trace plots
fit <- trim(fit, 8000, 2)

# Option 1 - calculate ALC from MCMC iterations
alc <- ALC(fit, xx)

# Option 2 - calculate ALC after predictions
fit <- predict(fit, xx)
alc <- ALC(fit)

# Visualize fit
plot(fit)
par(new = TRUE) # overlay ALC
plot(xx, alc$value, type = 'l', lty = 2, axes = FALSE, xlab = '', ylab = '')

# Select next design point
x_new <- xx[which.max(alc$value)]

# Evaluate fit
rmse(yy, fit$mean) # lower is better
score(yy, fit$mean, fit$Sigma) # higher is better

```

```

# 3. Three Layer and IMSPE -----
f <- function(x) {
  i <- which(x <= 0.48)
  x[i] <- 2 * sin(pi * x[i] * 4) + 0.4 * cos(pi * x[i] * 16)
  x[-i] <- 2 * x[-i] - 1
  return(x)
}

# Training data
x <- seq(0, 1, length = 30)
y <- f(x) + rnorm(30, 0, 0.05)

# Testing data
xx <- seq(0, 1, length = 100)
yy <- f(xx)

# Conduct MCMC
fit <- fit_three_layer(x, y, D = 1, nmcmc = 10000)
plot(fit) # investigate trace plots
fit <- trim(fit, 8000, 2)

# Option 1 - calculate IMSPE from only MCMC iterations
imspe <- IMSPE(fit, xx)

# Option 2 - calculate IMSPE after predictions
fit <- predict(fit, xx)
imspe <- IMSPE(fit)

# Visualize fit
plot(fit)
par(new = TRUE) # overlay IMSPE
plot(xx, imspe$value, type = 'l', lty = 2, axes = FALSE, xlab = '', ylab = '')

# Select next design point
x_new <- xx[which.min(imspe$value)]

# Evaluate fit
rmse(yy, fit$mean) # lower is better
score(yy, fit$mean, fit$Sigma) # higher is better

```

Description

Acts on a "gp", "dgp2", or "dgp3" object. Calculates ALC over the input locations x_{new} using specified reference grid. If no reference grid is specified, x_{new} is used as the reference. Optionally

utilizes SNOW parallelization. User should select the point with the highest ALC to add to the design.

Usage

```
ALC(object, x_new, ref, cores)

## S3 method for class 'gp'
ALC(object, x_new = NULL, ref = NULL, cores = 1)

## S3 method for class 'dgp2'
ALC(object, x_new = NULL, ref = NULL, cores = 1)

## S3 method for class 'dgp3'
ALC(object, x_new = NULL, ref = NULL, cores = 1)
```

Arguments

object	object from <code>fit_one_layer</code> , <code>fit_two_layer</code> , or <code>fit_three_layer</code>
x_new	matrix of possible input locations, if object has been run through <code>predict</code> the previously stored x_new is used
ref	optional reference grid for ALC approximation, if <code>ref = NULL</code> then x_new is used
cores	number of cores to utilize in parallel, by default no parallelization is used

Details

All iterations in the object are used in the calculation, so samples should be burned-in. Thinning the samples using `trim` will speed up computation. This function may be used in two ways:

- called on an object with only MCMC iterations, in which case `x_new` must be specified
- called on an object that has been predicted over, in which case the `x_new` from `predict` is used

SNOW parallelization reduces computation time but requires more memory storage.

C code derived from the "laGP" package (Robert B Gramacy and Furong Sun).

Value

list with elements:

- `value`: vector of ALC values, indices correspond to `x_new`
- `time`: computation time in seconds

References

Seo, S, M Wallat, T Graepel, and K Obermayer. 2000. "Gaussian Process Regression: Active Data Selection and Test Point Rejection." In *Mustererkennung 2000*, 27–34. New York, NY: Springer-Verlag.

Gramacy, RB and F Sun. (2016). laGP: Large-Scale Spatial Modeling via Local Approximate Gaussian Processes in R. *Journal of Statistical Software* 72 (1), 1-46. doi:10.18637/jss.v072.i01

Examples

```
# See "deepgp-package" or "fit_two_layer" for an example
```

calc_K	<i>Calculates covariance matrix</i>
--------	-------------------------------------

Description

Calculates covariance matrix based on inverse exponentiated squared euclidean distance with specified hyperparameters.

Usage

```
calc_K(d2, theta, g = NULL)
```

Arguments

d2	matrix of squared distances among input locations
theta	length scale parameter determining strength of correlation
g	nugget parameter determining noise level (only used if d2 is square)

Value

symmetric covariance matrix

Examples

```
x <- seq(0, 1, length = 10)
K <- calc_K(sq_dist(x), theta = 0.1, g = 0.01)
```

`continue`*Continues MCMC sampling*

Description

Acts on a "gp", "dgp2", or "dgp3" object. Continues MCMC sampling of hyperparameters and hidden layers and appends results to existing object.

Usage

```
continue(object, new_mcmc, trace)

## S3 method for class 'gp'
continue(object, new_mcmc = 1000, trace = TRUE)

## S3 method for class 'dgp2'
continue(object, new_mcmc = 1000, trace = TRUE)

## S3 method for class 'dgp3'
continue(object, new_mcmc = 1000, trace = TRUE)
```

Arguments

<code>object</code>	object from <code>fit_one_layer</code> , <code>fit_two_layer</code> , or <code>fit_three_layer</code>
<code>new_mcmc</code>	number of MCMC iterations to conduct and append
<code>trace</code>	logical indicating whether to print iteration progress

Details

See "fit_one_layer", "fit_two_layer", or "fit_three_layer" for details on MCMC. The resulting object will have `nmc` equal to the previous `nmc` plus `new_mcmc`. It is recommended to start an MCMC fit then investigate trace plots to assess burn-in. The primary use of this function is to gather more MCMC iterations in order to obtain burned-in samples.

Value

object of the same class with the new iterations appended

Examples

```
# See "deepgp-package" or "fit_two_layer" for an example
```


Description

Acts on a "gp", "dgp2", or "dgp3" object. Calculates expected improvement over input locations `x_new` with the goal of MINIMIZING the function. Optionally utilizes SNOW parallelization. User should select the point with the highest EI to add to the design.

Usage

```
EI(object, cores)

## S3 method for class 'gp'
EI(object, cores = 1)

## S3 method for class 'dgp2'
EI(object, cores = 1)

## S3 method for class 'dgp3'
EI(object, cores = 1)
```

Arguments

<code>object</code>	object from <code>fit_one_layer</code> , <code>fit_two_layer</code> , or <code>fit_three_layer</code> that has been predicted over <code>x_new</code> with <code>lite = FALSE</code>
<code>cores</code>	number of cores to utilize in parallel, by default no parallelization is used

Details

The object must be an output of `predict` with `lite = FALSE`. This will store the posterior mean and point-wise variance for every iteration. Once prediction is done, computation is relatively quick, so SNOW parallelization is only recommended when `nmc` is large.

Value

list with elements:

- `value`: vector of EI values, indices correspond to `x_new`
- `time`: computation time in seconds

References

Jones, DR, M Schonlau, and WJ Welch. 1998. "Efficient Global Optimization of Expensive Black-Box Functions." *Journal of Global Optimization* 13, 455-492. doi:10.1023/A:1008306431147.

Examples

```
# See "deepgp-package" or "fit_one_layer" for an example
```

fit_one_layer	<i>MCMC sampling for one layer GP</i>
---------------	---------------------------------------

Description

Conducts MCMC sampling of hyperparameters for a one layer GP. Covariance structure is based on inverse exponentiated squared euclidean distance with length scale parameter "theta" governing the strength of the correlation and nugget parameter "g" governing noise.

Usage

```
fit_one_layer(
  x,
  y,
  nmcmc = 10000,
  trace = TRUE,
  g_0 = 0.01,
  theta_0 = 0.5,
  true_g = NULL
)
```

Arguments

x	vector or matrix of input locations
y	vector of response values
nmcmc	number of MCMC iterations
trace	logical indicating whether to print iteration progress
g_0	initial value for g
theta_0	initial value for theta
true_g	if true nugget is known it may be specified here (set to a small value to make fit deterministic)

Details

Utilizes Metropolis Hastings sampling of the length scale and nugget parameters with a uniform proposal function (ranging from half to twice the previous iteration) and the following priors:

- `prior(g) <-dgamma(g, shape = 1.5, rate = 3.9)`
- `prior(theta) <-dgamma(theta, shape = 1.5, rate = 3.9/4)`

These priors are designed for "x" scaled to [0,1] and "y" scaled to have mean zero and variance 1. The output object of class "gp" is designed for use with `continue`, `trim`, and `predict`.

Value

a list of the S3 class "gp" with elements:

- x: copy of input matrix
- y: copy of response vector
- nmcmc: number of MCMC iterations
- g: vector of MCMC samples for g
- theta: vector of MCMC samples for theta
- time: computation time in seconds

References

Gramacy, RB. *Surrogates: Gaussian Process Modeling, Design, and Optimization for the Applied Sciences*. Chapman Hall, 2020.

Examples

```
# Toy example (runs in less than 5 seconds) -----
# This example uses a small number of MCMC iterations in order to run quickly
# More iterations are required to get appropriate fits
# Function defaults are recommended (see additional example below)

f <- function(x) {
  if (x <= 0.4) return(-1)
  if (x >= 0.6) return(1)
  if (x > 0.4 & x < 0.6) return(10*(x-0.5))
}
x <- seq(0.05, 0.95, length = 7)
y <- sapply(x, f)
x_new <- seq(0, 1, length = 100)

# Fit model and calculate EI
fit <- fit_one_layer(x, y, nmcmc = 500)
fit <- trim(fit, 400)
fit <- predict(fit, x_new, lite = FALSE)
ei <- EI(fit)

# One Layer and EI -----

f <- function(x) {
  sin(5 * pi * x) / (2 * x) + (x - 1) ^ 4
}

# Training data
x <- seq(0.5, 2, length = 30)
y <- f(x) + rnorm(30, 0, 0.01)

# Testing data
xx <- seq(0.5, 2, length = 100)
```

```

yy <- f(xx)

# Standardize inputs and outputs
xx <- (xx - min(x)) / (max(x) - min(x))
x <- (x - min(x)) / (max(x) - min(x))
yy <- (yy - mean(y)) / sd(y)
y <- (y - mean(y)) / sd(y)

# Conduct MCMC
fit <- fit_one_layer(x, y, nmcmc = 10000)
plot(fit) # investigate trace plots
fit <- trim(fit, 8000, 2)

# Predict and calculate EI
fit <- predict(fit, xx, lite = FALSE)
ei <- EI(fit)

# Visualize Fit
plot(fit)
par(new = TRUE) # overlay EI
plot(xx, ei$value, type = 'l', lty = 2, axes = FALSE, xlab = '', ylab = '')

# Select next design point
x_new <- xx[which.max(ei$value)]

# Evaluate fit
rmse(yy, fit$mean) # lower is better
score(yy, fit$mean, fit$Sigma) # higher is better

```

fit_three_layer

MCMC sampling for three layer deep GP

Description

Conducts MCMC sampling of hyperparameters, hidden layer "z", and hidden layer "w" for a three layer deep GP. Covariance structure is based on inverse exponentiated squared euclidean distance. Separate length scale parameters "theta_z", "theta_w", and "theta_y" govern the correlation strength of the inner layer, middle layer, and outer layer respectively. Nugget parameter "g" governs noise on the outer layer.

Usage

```

fit_three_layer(
  x,
  y,
  D = ifelse(is.matrix(x), ncol(x), 1),
  nmcmc = 10000,

```

```

trace = TRUE,
w_0 = suppressWarnings(matrix(x, nrow = length(y), ncol = D)),
z_0 = suppressWarnings(matrix(x, nrow = length(y), ncol = D)),
g_0 = 0.01,
theta_y_0 = 0.5,
theta_w_0 = 1,
theta_z_0 = 1,
true_g = NULL
)

```

Arguments

x	vector or matrix of input locations
y	vector of response values
D	integer designating dimension of hidden layers, defaults to dimension of x
nmcmc	number of MCMC iterations
trace	logical indicating whether to print iteration progress
w_0	initial value for hidden layer w, defaults to identity mapping (must be matrix of dimension nrow(x) by D or dimension nrow(x) - 1 by D)
z_0	initial value for hidden layer z, defaults to identity mapping (must be matrix of dimension nrow(x) by D or dimension nrow(x) - 1 by D)
g_0	initial value for g
theta_y_0	initial value for theta_y (length scale of outer layer)
theta_w_0	initial value for theta_w (length scale of middle layer), may be single value or vector of length D
theta_z_0	initial value for theta_z (length scale of inner layer), may be single value or vector of length D
true_g	if true nugget is known it may be specified here (set to a small value to make fit deterministic)

Details

Maps inputs "x" through hidden layer "z" then hidden layer "w" to outputs "y". Conducts sampling of the hidden layers using Elliptical Slice sampling. Utilizes Metropolis Hastings sampling of the length scale and nugget parameters with a uniform proposal function (ranging from half to twice the previous iteration) and the following priors:

- `prior(g) <- dgamma(g, 1.5, 3.9)`
- `prior(theta_z) <- dgamma(theta_z, 1.5, 3.9/4)`
- `prior(theta_w) <- dgamma(theta_w, 1.5, 3.9/4)`
- `prior(theta_y) <- dgamma(theta_y, 1.5, 3.9/6)`

These priors are designed for "x" scaled to [0,1] and "y" scaled to have mean zero and variance 1. The output object of class "dgp3" is designed for use with `continue`, `trim`, and `predict`. If `z_0` and `w_0` are of dimension `nrow(x) - 1` by `D`, the final rows are predicted using kriging. This is helpful in sequential design when adding a new input location and starting the MCMC at the place where the previous MCMC left off.

Value

a list of the S3 class "dgp3" with elements:

- x: copy of input matrix
- y: copy of response vector
- nmcmc: number of MCMC iterations
- g: vector of MCMC samples for g
- theta_y: vector of MCMC samples for theta_y (length scale of outer layer)
- theta_w: matrix of MCMC samples for theta_w (length scale of middle layer)
- theta_z: matrix of MCMC samples for theta_z (length scale of inner layer)
- w: list of MCMC samples for middle hidden layer w
- z: list of MCMC samples for inner hidden layer z
- time: computation time in seconds

References

Damianou, A and N Lawrence. (2013). "Deep gaussian processes." *Artificial Intelligence and Statistics*, 207-215.

Murray, I, RP Adams, and D MacKay. 2010. "Elliptical slice sampling." *Journal of Machine Learning Research* 9, 541-548.

Examples

```
# Toy example (runs in less than 5 seconds) -----
# This example uses a small number of MCMC iterations in order to run quickly
# More iterations are required to get appropriate fits
# Function defaults are recommended (see additional example below)

f <- function(x) {
  if (x <= 0.4) return(-1)
  if (x >= 0.6) return(1)
  if (x > 0.4 & x < 0.6) return(10*(x-0.5))
}
x <- seq(0.05, 0.95, length = 7)
y <- sapply(x, f)
x_new <- seq(0, 1, length = 100)

# Fit model and calculate IMSPE
fit <- fit_three_layer(x, y, nmcmc = 500)
fit <- trim(fit, 400)
fit <- predict(fit, x_new)
imspe <- IMSPE(fit)

# Three Layer and IMSPE -----

f <- function(x) {
```

```

    i <- which(x <= 0.48)
    x[i] <- 2 * sin(pi * x[i] * 4) + 0.4 * cos(pi * x[i] * 16)
    x[-i] <- 2 * x[-i] - 1
    return(x)
}

# Training data
x <- seq(0, 1, length = 30)
y <- f(x) + rnorm(30, 0, 0.05)

# Testing data
xx <- seq(0, 1, length = 100)
yy <- f(xx)

# Conduct MCMC
fit <- fit_three_layer(x, y, D = 1, nmcmc = 10000)
plot(fit) # investigate trace plots
fit <- trim(fit, 8000, 2)

# Option 1 - calculate IMSPE from only MCMC iterations
imspe <- IMSPE(fit, xx)

# Option 2 - calculate IMSPE after predictions
fit <- predict(fit, xx)
imspe <- IMSPE(fit)

# Visualize fit
plot(fit)
par(new = TRUE) # overlay IMSPE
plot(xx, imspe$value, type = 'l', lty = 2, axes = FALSE, xlab = '', ylab = '')

# Select next design point
x_new <- xx[which.min(imspe$value)]

# Evaluate fit
rmse(yy, fit$mean) # lower is better
score(yy, fit$mean, fit$Sigma) # higher is better

```

fit_two_layer

MCMC sampling for two layer deep GP

Description

Conducts MCMC sampling of hyperparameters and hidden layer "w" for a two layer deep GP. Covariance structure is based on inverse exponentiated squared euclidean distance. Separate length scale parameters "theta_w" and "theta_y" govern the correlation strength of the hidden layer and outer layer respectively. Nugget parameter "g" governs noise on the outer layer.

Usage

```
fit_two_layer(
  x,
  y,
  D = ifelse(is.matrix(x), ncol(x), 1),
  nmcmc = 10000,
  trace = TRUE,
  w_0 = suppressWarnings(matrix(x, nrow = length(y), ncol = D)),
  g_0 = 0.01,
  theta_y_0 = 0.5,
  theta_w_0 = 1,
  true_g = NULL
)
```

Arguments

x	vector or matrix of input locations
y	vector of response values
D	integer designating dimension of hidden layer, defaults to dimension of x
nmcmc	number of MCMC iterations
trace	logical indicating whether to print iteration progress
w_0	initial value for hidden layer w, defaults to identity mapping (must be matrix of dimension nrow(x) by D or dimension nrow(x) -1 by D)
g_0	initial value for g
theta_y_0	initial value for theta_y (length scale of outer layer)
theta_w_0	initial value for theta_w (length scale of inner layer), may be single value or vector of length D
true_g	if true nugget is known it may be specified here (set to a small value to make fit deterministic)

Details

Maps inputs "x" through hidden layer "w" to outputs "y". Conducts sampling of the hidden layer using Elliptical Slice sampling. Utilizes Metropolis Hastings sampling of the length scale and nugget parameters with a uniform proposal function (ranging from half to twice the previous iteration) and the following priors:

- `prior(g) <- dgamma(g, 1.5, 3.9)`
- `prior(theta_w) <- dgamma(theta_w, 1.5, 3.9/4)`
- `prior(theta_y) <- dgamma(theta_y, 1.5, 3.9/6)`

These priors are designed for "x" scaled to [0,1] and "y" scaled to have mean zero and variance 1. The output object of class "dgp2" is designed for use with `continue`, `trim`, and `predict`. If `w_0` is of dimension `nrow(x) - 1` by `D`, the final row is predicted using kriging. This is helpful in sequential design when adding a new input location and starting the MCMC at the place where the previous MCMC left off.

Value

a list of the S3 class "dgp2" with elements:

- x: copy of input matrix
- y: copy of response vector
- nmcmc: number of MCMC iterations
- g: vector of MCMC samples for g
- theta_y: vector of MCMC samples for theta_y (length scale of outer layer)
- theta_w: matrix of MCMC samples for theta_w (length scale of inner layer)
- w: list of MCMC samples for hidden layer w
- time: computation time in seconds

References

Damianou, A and N Lawrence. (2013). "Deep gaussian processes." *Artificial Intelligence and Statistics*, 207-215.

Murray, I, RP Adams, and D MacKay. 2010. "Elliptical slice sampling." *Journal of Machine Learning Research* 9, 541-548.

Examples

```
# Toy example (runs in less than 5 seconds) -----
# This example uses a small number of MCMC iterations in order to run quickly
# More iterations are required to get appropriate fits
# Function defaults are recommended (see additional example below)

f <- function(x) {
  if (x <= 0.4) return(-1)
  if (x >= 0.6) return(1)
  if (x > 0.4 & x < 0.6) return(10*(x-0.5))
}
x <- seq(0.05, 0.95, length = 7)
y <- sapply(x, f)
x_new <- seq(0, 1, length = 100)

# Fit model and calculate ALC
fit <- fit_two_layer(x, y, nmcmc = 500)
fit <- trim(fit, 400)
fit <- predict(fit, x_new)
alc <- ALC(fit)

# Two Layer and ALC -----

f <- function(x) {
  exp(-10 * x) * (cos(10 * pi * x - 1) + sin(10 * pi * x - 1)) * 5 - 0.2
}
```

```

# Training data
x <- seq(0, 1, length = 30)
y <- f(x) + rnorm(30, 0, 0.05)

# Testing data
xx <- seq(0, 1, length = 100)
yy <- f(xx)

# Conduct MCMC
fit <- fit_two_layer(x, y, D = 1, nmcnc = 9000)
fit <- continue(fit, 1000)
plot(fit) # investigate trace plots
fit <- trim(fit, 8000, 2)

# Option 1 - calculate ALC from MCMC iterations
alc <- ALC(fit, xx)

# Option 2 - calculate ALC after predictions
fit <- predict(fit, xx)
alc <- ALC(fit)

# Visualize fit
plot(fit)
par(new = TRUE) # overlay ALC
plot(xx, alc$value, type = 'l', lty = 2, axes = FALSE, xlab = '', ylab = '')

# Select next design point
x_new <- xx[which.max(alc$value)]

# Evaluate fit
rmse(yy, fit$mean) # lower is better
score(yy, fit$mean, fit$Sigma) # higher is better

```

Description

Acts on a "gp", "dgp2", or "dgp3" object. Calculates IMSPE over the input locations `x_new`. Optionally utilizes SNOW parallelization. User should select the point with the lowest IMSPE to add to the design.

Usage

```

IMSPE(object, x_new, cores)

## S3 method for class 'gp'
IMSPE(object, x_new = NULL, cores = 1)

```

```
## S3 method for class 'dgp2'  
IMSPE(object, x_new = NULL, cores = 1)  
  
## S3 method for class 'dgp3'  
IMSPE(object, x_new = NULL, cores = 1)
```

Arguments

object	object from <code>fit_one_layer</code> , <code>fit_two_layer</code> , or <code>fit_three_layer</code>
x_new	matrix of possible input locations, if object has been run through <code>predict</code> the previously stored x_new is used
cores	number of cores to utilize in parallel, by default no parallelization is used

Details

All iterations in the object are used in the calculation, so samples should be burned-in. Thinning the samples using `trim` will speed up computation. This function may be used in two ways:

- called on an object with only MCMC iterations, in which case `x_new` must be specified
- called on an object that has been predicted over, in which case the `x_new` from `predict` is used

SNOW parallelization reduces computation time but requires more memory storage.

Value

list with elements:

- `value`: vector of IMSPE values, indices correspond to `x_new`
- `time`: computation time in seconds

References

Binois, M, J Huang, RB Gramacy, and M Ludkovski. 2019. "Replication or Exploration? Sequential Design for Stochastic Simulation Experiments." *Technometrics* 61, 7-23. Taylor & Francis. doi:10.1080/00401706.2018.1469433.

Examples

```
# See "deepgp-package" or "fit_three_layer" for an example
```

plot

Plots object from "deepgp" package

Description

Acts on a "gp", "dgp2", or "dgp3" object. Generates trace plots for length scale and nugget hyper-parameters. Generates plots of hidden layers for one-dimensional inputs. Generates plots of the posterior mean and estimated 95% prediction intervals for one-dimensional inputs; generates heat maps of the posterior mean and point-wise variance for two-dimensional inputs.

Usage

```
## S3 method for class 'gp'  
plot(x, trace = TRUE, predict = TRUE, ...)  
  
## S3 method for class 'dgp2'  
plot(x, trace = TRUE, hidden = FALSE, predict = TRUE, ...)  
  
## S3 method for class 'dgp3'  
plot(x, trace = TRUE, hidden = FALSE, predict = TRUE, ...)
```

Arguments

x	object from <code>fit_one_layer</code> , <code>fit_two_layer</code> , or <code>fit_three_layer</code>
trace	logical indicating whether to generate trace plots
predict	logical indicating whether to generate posterior predictive plot
...	N/A
hidden	logical indicating whether to generate plots of hidden layers ("dgp2" or "dgp3" only)

Details

Trace plots are useful in assessing burn-in. Hidden layer plots are colored on a gradient - red lines represent earlier iterations and yellow lines represent later iterations - to help assess burn-in of the hidden layers. These plots are meant to help in model fitting and visualization.

Examples

```
# See "deepgp-package", "fit_one_layer", "fit_two_layer", or "fit_three_layer"  
# for an example
```

predict	<i>Predict posterior mean and covariance</i>
---------	--

Description

Acts on a "gp", "dgp2", or "dgp3" object. Calculates posterior mean and covariance over specified input locations. Optionally utilizes SNOW parallelization.

Usage

```
## S3 method for class 'gp'
predict(object, x_new, lite = TRUE, cores = 1, ...)

## S3 method for class 'dgp2'
predict(
  object,
  x_new,
  lite = TRUE,
  cores = 1,
  uncertainty = c("CONDITIONAL", "FULL"),
  ...
)

## S3 method for class 'dgp3'
predict(
  object,
  x_new,
  lite = TRUE,
  cores = 1,
  uncertainty = c("CONDITIONAL", "FULL"),
  ...
)
```

Arguments

object	object from <code>fit_one_layer</code> , <code>fit_two_layer</code> , or <code>fit_three_layer</code> with burn-in already removed
x_new	matrix of predictive input locations
lite	logical indicating whether to store the mean and diagonal of the covariance for every iteration (must use <code>lite = FALSE</code> in order to use EI)
cores	number of cores to utilize in parallel, by default no parallelization is used
...	N/A
uncertainty	denotes whether to incorporate conditional or full predictive uncertainty in mapping through hidden layers ("dgp2" or "dgp3" only)

Details

All iterations in the object are used for prediction, so samples should be burned-in. Thinning the samples using `trim` will speed up computation. The posterior mean and covariance are calculated for each iteration, then averaged. The covariance of the means is appropriately added to the average of the covariances.

SNOW parallelization reduces computation time but requires significantly more memory storage. Use `cores = 1` if memory is limited.

Value

object of the same class with the following additional elements:

- `x_new`: copy of predictive input locations
- `tau2`: vector of `tau2` estimates (governing the magnitude of the covariance)
- `mean`: predicted posterior mean, indices correspond to `x_new` location
- `Sigma`: predicted posterior covariance, indices correspond to `x_new` location
- `Sigma_smooth`: predicted posterior covariance with `g` removed from the diagonal
- `mu_t`: (only when `lite = FALSE`) matrix of posterior mean for each iteration, column index corresponds to iteration and row index corresponds to `x_new` location
- `sig2_t`: (only when `lite = FALSE`) matrix of posterior point-wise variance (diagonal of `Sigma`) for each iteration, column index corresponds to iteration and row index corresponds to `x_new` location
- `w_new`: ("dgp2" and "dgp3" only) list of hidden layer predictions, list index corresponds to iteration and row index corresponds to `x_new` location
- `z_new`: ("dgp3" only) list of hidden layer predictions, list index corresponds to iteration and row index corresponds to `x_new` location

Computation time is added to the computation time of the existing object.

References

Gramacy, RB. *Surrogates: Gaussian Process Modeling, Design, and Optimization for the Applied Sciences*. Chapman Hall, 2020.

Examples

```
# See "deepgp-package", "fit_one_layer", "fit_two_layer", or "fit_three_layer"
# for an example
```

rmse	<i>Calculates RMSE</i>
------	------------------------

Description

Calculates root mean square error (lower RMSE indicate better fits).

Usage

```
rmse(y, mu)
```

Arguments

y	response vector
mu	predicted mean

Examples

```
# See "deepgp-package", "fit_one_layer", "fit_two_layer", or "fit_three_layer"  
# for an example
```

score	<i>Calculates score</i>
-------	-------------------------

Description

Calculates score (higher scores indicate better fits). Only applicable to noisy data.

Usage

```
score(y, mu, sigma)
```

Arguments

y	response vector
mu	predicted mean
sigma	predicted covariance

References

Gneiting, T, and AE Raftery. 2007. "Strictly Proper Scoring Rules, Prediction, and Estimation." *Journal of the American Statistical Association* 102 (477), 359-378.

Examples

```
# See "deepgp-package", "fit_one_layer", "fit_two_layer", or "fit_three_layer"  
# for an example
```

sq_dist	<i>Calculates squared pairwise distances</i>
---------	--

Description

Calculates squared pairwise euclidean distances using C.

Usage

```
sq_dist(X1, X2 = NULL)
```

Arguments

X1	matrix of input locations
X2	matrix of second input locations (if NULL, distance is calculated between X1 and itself)

Details

C code derived from the "laGP" package (Robert B Gramacy and Furong Sun).

Value

symmetric matrix of squared euclidean distances

References

Gramacy, RB and F Sun. (2016). laGP: Large-Scale Spatial Modeling via Local Approximate Gaussian Processes in R. *Journal of Statistical Software* 72 (1), 1-46. doi:10.18637/jss.v072.i01

Examples

```
x <- seq(0, 1, length = 10)  
d2 <- sq_dist(x)
```

trim	<i>Trim/Thin MCMC iterations</i>
------	----------------------------------

Description

Acts on a "gp", "dgp2", or "dgp3" object. Removes the specified number of MCMC iterations (starting at the first iteration). After these samples are removed, the remaining samples may be thinned.

Usage

```
trim(object, burn, thin)

## S3 method for class 'gp'
trim(object, burn, thin = 1)

## S3 method for class 'dgp2'
trim(object, burn, thin = 1)

## S3 method for class 'dgp3'
trim(object, burn, thin = 1)
```

Arguments

object	object from <code>fit_one_layer</code> , <code>fit_two_layer</code> , or <code>fit_three_layer</code>
burn	integer specifying number of iterations to cut off as burn-in
thin	integer specifying amount of thinning (thin = 1 keeps all iterations, thin = 2 keeps every other iteration, thin = 10 keeps every tenth iteration, etc.)

Details

The resulting object will have `nmc` equal to the previous `nmc` minus `burn` divided by `thin`. It is recommended to start an MCMC fit then investigate trace plots to assess burn-in. Once burn-in has been achieved, use this function to remove the starting iterations. Thinning reduces the size of the resulting object while accounting for the high correlation between consecutive iterations.

Value

object of the same class with the selected iterations removed

Examples

```
# See "deepgp-package", "fit_one_layer", "fit_two_layer", or "fit_three_layer"
# for an example
```

Index

ALC, [2](#), [5](#)

calc_K, [7](#)

continue, [2](#), [8](#)

deepgp-package, [2](#)

EI, [2](#), [9](#)

fit_one_layer, [2](#), [10](#)

fit_three_layer, [2](#), [12](#)

fit_two_layer, [2](#), [15](#)

IMSPE, [2](#), [18](#)

plot, [2](#), [20](#)

predict, [2](#), [21](#)

rmse, [23](#)

score, [23](#)

sq_dist, [24](#)

trim, [2](#), [25](#)