

Package: FastKRR (via r-universe)

June 13, 2026

Type Package

Title Kernel Ridge Regression using 'RcppArmadillo'

Version 0.1.2

Description Provides core computational operations in C++ via 'RcppArmadillo', enabling faster performance than pure R, improved numerical stability, and parallel execution with OpenMP where available. On systems without OpenMP support, the package automatically falls back to single-threaded execution with no user configuration required. For efficient model selection, it integrates with 'CVST' to provide sequential-testing cross-validation that identifies competitive hyperparameters without exhaustive grid search. The package offers a unified interface for exact kernel ridge regression and three scalable approximations—Nyström, Pivoted Cholesky, and Random Fourier Features—allowing analyses with substantially larger sample sizes than are feasible with exact KRR. It also integrates with the 'tidymodels' ecosystem via the 'parsnip' model specification 'krr_reg', and the S3 method `tunable.krr_reg()`. To understand the theoretical background, one can refer to Wainwright (2019) <[doi:10.1017/9781108627771](https://doi.org/10.1017/9781108627771)>.

License GPL (>= 2)

URL <https://github.com/kybak90/FastKRR>, <https://www.tidymodels.org>

BugReports <https://github.com/kybak90/FastKRR/issues>

Imports CVST, generics, parsnip, Rcpp, rlang, tibble, ggplot2

LinkingTo Rcpp, RcppArmadillo

Suggests knitr, rmarkdown, dials, tidymodels, modeldata, dplyr

SystemRequirements OpenMP (optional)

Encoding UTF-8

RoxygenNote 7.3.3

NeedsCompilation yes

Author Gyeongmin Kim [aut] (Sungshin Women's University), Seyoung Lee [aut] (Sungshin Women's University), Miyoung Jang [aut] (Sungshin Women's University), Kwan-Young Bak [aut, cre, cph] (ORCID: <<https://orcid.org/0000-0002-4541-160X>>, Sungshin Women's University)

Maintainer Kwan-Young Bak <kybak@sungshin.ac.kr>

Repository <https://cran.r-universe.dev>

Date/Publication 2025-11-15 11:20:02 UTC

RemoteUrl <https://github.com/cran/FastKRR>

RemoteRef HEAD

RemoteSha 413db09125a766a2a3f8c80069a7c219cc380e2d

Contents

FastKRR-package	2
approx_kernel	3
coef.krr	6
error	7
error.krr	8
fastkrr	9
krr_reg	12
make_kernel	15
param	16
param.krr	16
plot.krr	18
predict.krr	19
print.approx_kernel	20
print.kernel_matrix	21
print.krr	22
summary.krr	23
tunable.krr_reg	24
Index	25

FastKRR-package

*Kernel Ridge Regression using the **RcppArmadillo** Package*

Description

The **FastKRR** implements its core computational operations in C++ via **RcppArmadillo**, enabling faster performance than pure R, improved numerical stability, and parallel execution with OpenMP where available. On systems without OpenMP support, the package automatically falls back to single-threaded execution with no user configuration required. For efficient model selection, it integrates with **CVST** to provide sequential-testing cross-validation that identifies competitive hyperparameters without exhaustive grid search. The package offers a unified interface for exact kernel

ridge regression and three widely used scalable approximations—Nyström, Pivoted Cholesky, and Random Fourier Features—allowing analyses with substantially larger sample sizes than are feasible with exact KRR while retaining strong predictive performance. This combination of a compiled backend and scalable algorithms addresses limitations of packages that rely solely on exact computation, which is often impractical for large n . It also integrates with the **tidymodels** ecosystem via the **parsnip** model specification `krr_reg`, and the S3 method `tunable.krr_reg()` (exposes tunable parameters to `dials/tune`); see their help pages for usage.

Directory structure

- R/: High-level R functions and user-facing API
- src/: C++ sources (kernel computation, fitting, prediction)

This package links against **Rcpp** and **RcppArmadillo** (via `LinkingTo`). It uses **CVST**, **parsnip**, and the **tidymodels** ecosystem through their public R APIs.

Author(s)

Maintainer: Kwan-Young Bak <kybak@sunghsin.ac.kr> ([ORCID](#)) (Sunghsin Women’s University) [copyright holder]

Authors:

- Gyeongmin Kim <rlarudals0824@gmail.com> (Sunghsin Women’s University)
- Seyoung Lee <sudang0404@gmail.com> (Sunghsin Women’s University)
- Miyoung Jang <miyoung9072@gmail.com> (Sunghsin Women’s University)

See Also

CVST, **Rcpp**, **RcppArmadillo**, **parsnip**, **tidymodels**

approx_kernel

Compute low-rank approximations(Nyström, Pivoted Cholesky, RFF)

Description

Computes low-rank kernel approximation $\tilde{K} \in \mathbb{R}^{n \times n}$ using three methods: Nyström approximation, Pivoted Cholesky decomposition, and Random Fourier Features (RFF).

Usage

```
approx_kernel(
  K = NULL,
  X = NULL,
  opt = c("nystrom", "pivoted", "rff"),
  kernel = c("gaussian", "laplace"),
  m = NULL,
  d,
```

```

rho,
eps = 1e-06,
W = NULL,
b = NULL,
n_threads = 4
)

```

Arguments

K	Exact Kernel matrix $K \in \mathbb{R}^{n \times n}$. Used in "nystrom" and "pivoted".
X	Design matrix $X \in \mathbb{R}^{n \times d}$. Only required for "rff".
opt	Method for constructing or approximating : "nystrom" Construct a low-rank approximation of the kernel matrix $K \in \mathbb{R}^{n \times n}$ using the Nyström approximation. "pivoted" Construct a low-rank approximation of the kernel matrix $K \in \mathbb{R}^{n \times n}$ using Pivoted Cholesky decomposition. "rff" Construct a low-rank approximation of the kernel matrix $K \in \mathbb{R}^{n \times n}$ using Random Fourier Features (RFF).
kernel	Kernel type either "gaussian" or "laplace".
m	Approximation rank (number of random features) for the low-rank kernel approximation. If not specified, the recommended choice is $\lceil n \cdot \log(d + 5) / 10 \rceil$ where X is design matrix, $n = \text{row}(X)$ and $d = \text{ncol}(X)$.
d	Design matrix's dimension ($d = \text{ncol}(X)$).
rho	Scaling parameter of the kernel (ρ), specified by the user.
eps	Tolerance parameter used only in "pivoted" for stopping criterion of the Pivoted Cholesky decomposition.
W	Random frequency matrix $\omega \in \mathbb{R}^{m \times d}$
b	Random phase vector $b \in \mathbb{R}^m$, i.i.d. $\text{Unif}[0, 2\pi]$.
n_threads	Number of parallel threads. The default is 4. If the system does not support 4 threads, it automatically falls back to 1 thread. It is applied only for opt = "nystrom" or opt = "rff", and for the Laplace kernel (kernel = "laplace").

Details

Requirements and what to supply:

Common

- d and rho must be provided (non-NULL).

nystrom / pivoted

- Require a precomputed kernel matrix K; error if K is NULL.
- If m is NULL, use $\lceil n \cdot \log(d + 5) / 10 \rceil$.

- For "pivoted", a tolerance `eps` is used; the decomposition stops early when the next pivot (residual diagonal) drops below `eps`.

rff

- `K` must be `NULL` (not used) and `X` must be provided with `d = ncol(X)`.
- The function automatically generates `W` (random frequency matrix $\omega \in \mathbb{R}^{m \times d}$) and `b` (random phase vector $b \in \mathbb{R}^m$).
- If the user provides them manually, both `W` and `b` must be specified and their dimensions must be compatible.

Value

An S3 object of class "approx_kernel" containing the results of the kernel approximation:

- `call`: The matched function call used to create the object.
- `opt`: The kernel approximation method actually used ("nystrom", "pivoted", "rff").
- `K_approx`: $n \times n$ approximated kernel matrix.
- `m`: Kernel approximation degree.

Additional components depend on the value of `opt`:

nystrom

- `n_threads`: Number of threads used in the computation.

pivoted

- `eps`: Numerical tolerance used for early stopping in the pivoted Cholesky decomposition.

rff

- `d`: Input design matrix's dimension.
- `rho`: Scaling parameter of the kernel.
- `W`: $m \times d$ Random frequency matrix.
- `b`: Random phase m -vector.
- `used_supplied_Wb`: Logical; TRUE if user-supplied `W`, `b` were used, FALSE otherwise.
- `n_threads`: Number of threads used in the computation.

Examples

```
# Data setting
set.seed(1)
d = 1
n = 1000
m = 50
X = matrix(runif(n*d, 0, 1), nrow = n, ncol = d)
y = as.vector(sin(2*pi*rowMeans(X)^3) + rnorm(n, 0, 0.1))
K = make_kernel(X, kernel = "gaussian", rho = 1)
```

```

# Example: RFF approximation
K_rff = approx_kernel(X = X, opt = "rff", kernel = "gaussian",
                      m = m, d = d, rho = 1,
                      n_threads = 1)

# Exapmle: Nystrom approximation
K_nystrom = approx_kernel(K = K, opt = "nystrom",
                          m = m, d = d, rho = 1,
                          n_threads = 1)

# Example: Pivoted Cholesky approximation
K_pivoted = approx_kernel(K = K, opt = "pivoted",
                          m = m, d = d, rho = 1)

```

coef.krr

Coef method for fitted Kernel Ridge Regression models

Description

Displays the main coefficient information from a fitted Kernel Ridge Regression (KRR) model, including the original function call and the first few estimated coefficients. The type of coefficient reported depends on the kernel approximation method: for `opt = "exact"`, `"nystrom"`, or `"pivoted"`, the coefficients represent α ; for `opt = "rff"`, they represent the coefficient (β).

Usage

```

## S3 method for class 'krr'
coef(object, ...)

```

Arguments

object	An S3 object of class <code>krr</code> , typically returned by fastkrr .
...	Additional arguments (currently ignored).

Value

A human-readable summary of the fitted KRR model to the console.

See Also

[fastkrr](#)

Examples

```
# Data setting
set.seed(1)
lambda = 1e-4
d = 1
n = 50
rho = 1
X = matrix(runif(n*d, 0, 1), nrow = n, ncol = d)
y = as.vector(sin(2*pi*rowMeans(X)^3) + rnorm(n, 0, 0.1))

# Example: exact
model = fastkrr(X, y,
               kernel = "gaussian", opt = "exact",
               rho = rho, lambda = 1e-4)

class(model)

coef(model)
```

error

Compute Model Error (Generic)

Description

Generic function for computing model error.

Usage

```
error(x, ...)
```

```
## Default S3 method:
```

```
error(x, ...)
```

Arguments

x An object to compute model error for.

... Additional arguments passed to methods.

Value

A numeric value or class-specific result.

 error.krr

Compute Model Error for Kernel Ridge Regression Models

Description

Computes the model error for kernel ridge regression ("krr") objects. Returns the mean squared error (MSE) between the observed responses and the fitted values stored in the object.

Usage

```
## S3 method for class 'krr'
error(x, ...)
```

Arguments

x An object of class "krr", typically returned by [fastkrr](#).
 ... Additional arguments (ignored).

Details

This method computes:

$$\text{MSE} = \frac{1}{n} \sum_i (y_i - \hat{y}_i)^2$$

where `y` and `fitted.values` are stored in the "krr" object attributes.

Value

A numeric value giving the mean squared error (MSE).

See Also

[summary.krr](#), [plot.krr](#), [predict.krr](#)

Examples

```
# Data setting
set.seed(1)
lambda = 1e-4
d = 1
n = 50
rho = 1
X = matrix(runif(n*d, 0, 1), nrow = n, ncol = d)
y = as.vector(sin(2*pi*rowMeans(X)^3) + rnorm(n, 0, 0.1))

model = fastkrr(X, y, kernel = "gaussian", lambda = 0.001)
error(model)
```

fastkrr

*Fit kernel ridge regression using exact or approximate methods***Description**

This function performs kernel ridge regression (KRR) in high-dimensional settings. The regularization parameter λ can be selected via the CVST (Cross-Validation via Sequential Testing) procedure. For scalability, three different kernel approximation strategies are supported (Nyström approximation, Pivoted Cholesky decomposition, Random Fourier Features(RFF)), and kernel matrix can be computed using two methods(Gaussian kernel, Laplace kernel).

Usage

```
fastkrr(
  x,
  y,
  kernel = "gaussian",
  opt = "exact",
  m = NULL,
  eps = 1e-06,
  rho = 1,
  lambda = NULL,
  fastcv = FALSE,
  n_threads = 4,
  verbose = TRUE
)
```

Arguments

x	Design matrix $X \in \mathbb{R}^{n \times d}$.
y	Response variable $y \in \mathbb{R}^n$.
kernel	Kernel type either "gaussian" or "laplace".
opt	Method for constructing or approximating : "exact" Construct the full kernel matrix $K \in \mathbb{R}^{n \times n}$ using design matrix X . "nyström" Construct a low-rank approximation of the kernel matrix $K \in \mathbb{R}^{n \times n}$ using the Nyström approximation. "pivoted" Construct a low-rank approximation of the kernel matrix $K \in \mathbb{R}^{n \times n}$ using Pivoted Cholesky decomposition. "rff" Use Random Fourier Features to construct a feature map $Z \in \mathbb{R}^{n \times m}$ (with m random features) so that $K \approx ZZ^\top$. Here, m is the number of features.
m	Approximation rank(number of random features) used for the low-rank kernel approximation. If not provided by the user, it defaults to

$$\lceil n \cdot \frac{\log(d+5)}{10} \rceil,$$

where $n = nrow(X)$ and $d = ncol(X)$.

eps	Tolerance parameter used only in "pivoted" for stopping criterion of the Pivoted Cholesky decomposition.
rho	Scaling parameter of the kernel(ρ), specified by the user. Defaults to 1.
	Gaussian kernel : $\mathcal{K}(x, x') = \exp(-\rho\ x - x'\ _2^2)$
	Laplace kernel : $\mathcal{K}(x, x') = \exp(-\rho\ x - x'\ _1)$
lambda	Regularization parameter. If NULL, the penalty parameter is chosen automatically via CVST package. If not provided, the argument is set to a kernel-specific grid of 100 values: $[10^{-10}, 10^{-3}]$ for Gaussian, $[10^{-5}, 10^{-2}]$ for Laplace.
fastcv	If TRUE, accelerated cross-validation is performed via sequential testing (early stopping) as implemented in the CVST package. The default is FALSE.
n_threads	Number of parallel threads. The default is 4. If the system does not support 4 threads, it automatically falls back to 1 thread. Parallelization (implemented in C++) is one of the main advantages of this package and is applied only for <code>opt = "nystrom"</code> or <code>opt = "rff"</code> , and for the Laplace kernel (<code>kernel = "laplace"</code>).
verbose	If TRUE, detailed progress and cross-validation results are printed to the console. If FALSE, suppresses intermediate output and only returns the final result.

Details

The function performs several input checks and automatic adjustments:

- If `x` is a vector, it is converted to a one column matrix. Otherwise, `x` must be a matrix; otherwise an error is thrown.
- `y` must be a vector, and its length must match `nrow(x)`.
- `kernel` must be either `gaussian` or `laplace`.
- `opt` must be one of `"exact"`, `"pivoted"`, `"nystrom"`, or `"rff"`.
- If `m` is NULL, it defaults to

$$\lceil n \cdot \log(d + 5) / 10 \rceil$$

where $n = nrow(X)$ and $d = ncol(X)$. Otherwise, `m` must be a positive integer.

- `rho` must be a positive real number (default is 1).
- `lambda` can be specified in three ways:
 1. A positive numeric scalar, in which case the model is fitted with this single value.
 2. A numeric vector (length ≥ 3) of positive values used as a tuning grid; selection is performed by **CVST** cross-validation (sequential testing if `fastcv = TRUE`).
 3. NULL: use a default grid (internal setting) and tune `lambda` via **CVST** cross-validation (sequential testing if `fastcv = TRUE`).
- `n_threads`: Number of threads for parallel computation. Default is 4. If the system has ≤ 3 available processors, it uses 1.

Value

An S3 object of class "fastkrr", which is a list containing the results of the fitted Kernel Ridge Regression model.

- `coefficients`: Estimated coefficient vector \mathbb{R}^n . Accessible via `model$coefficients`.
- `fitted.values`: Fitted values \mathbb{R}^n . Accessible via `model$fitted.values`.
- `opt`: Kernel approximation option. One of "exact", "pivoted", "nystrom", "rff".
- `kernel`: Kernel used ("gaussian" or "laplace").
- `x`: Input design matrix.
- `y`: Response vector.
- `lambda`: Regularization parameter. If NULL, tuned by cross-validation via **CVST**.
- `rho`: Additional user-specified hyperparameter.
- `n_threads`: Number of threads used for parallelization.

Additional components depend on the value of `opt`:

opt = "exact":

- `K`: The full kernel matrix.

opt = "nystrom":

- `K`: Exact kernel matrix $K \in \mathbb{R}^{n \times n}$.
- `m`: Kernel approximation degree.
- `R`: The method provides a low-rank approximation to the kernel matrix $R \in \mathbb{R}^{n \times m}$ obtained via Nyström approximation; satisfies $K \approx RR^\top$.

opt = "pivoted":

- `K`: Exact kernel matrix $K \in \mathbb{R}^{n \times n}$.
- `m`: Kernel approximation degree.
- `PR`: The method provides a low-rank approximation to the kernel matrix $PR \in \mathbb{R}^{n \times m}$ obtained via Pivoted Cholesky decomposition; satisfies $K \approx PR(PR)^\top$.
- `eps`: Numerical tolerance used for early stopping in the pivoted Cholesky decomposition.

opt = "rff":

- `m`: Number of random features.
- `Z`: Random Fourier Feature matrix $Z \in \mathbb{R}^{n \times m}$ with $Z_{ij} = z_j(x_i) = \sqrt{2/m} \cos(\omega_j^\top x_i + b_j)$, $j = 1, \dots, m$, so that $K \approx ZZ^\top$.
- `W`: Random frequency matrix $\omega \in \mathbb{R}^{m \times d}$ (row j is $\omega_j^\top \in \mathbb{R}^d$), drawn i.i.d. from the spectral density of the chosen kernel:
 - Gaussian: $\omega_{jk} \sim \mathcal{N}(0, 2\gamma)$ (e.g., $\gamma = 1/\ell^2$).
 - Laplace: $\omega_{jk} \sim \text{Cauchy}(0, 1/\sigma)$ i.i.d.
- `b` Random phase vector $b \in \mathbb{R}^m$, i.i.d. $\text{Unif}[0, 2\pi]$.

Examples

```
# Data setting
set.seed(1)
lambda = 1e-4
d = 1
rho = 1
n = 50
X = matrix(runif(n*d, 0, 1), nrow = n, ncol = d)
y = as.vector(sin(2*pi*rowMeans(X)^3) + rnorm(n, 0, 0.1))

# Example: pivoted cholesky
model = fastkrr(X, y, kernel = "gaussian", opt = "pivoted", rho = rho, lambda = 1e-4)

# Example: nystrom
model = fastkrr(X, y, kernel = "gaussian", opt = "nystrom", rho = rho, lambda = 1e-4)

# Example: random fourier features
model = fastkrr(X, y, kernel = "gaussian", opt = "rff", rho = rho, lambda = 1e-4)

# Example: Laplace kernel
model = fastkrr(X, y, kernel = "laplace", opt = "nystrom", n_threads = 1, rho = rho)
```

krr_reg

Kernel Ridge Regression

Description

Defines a Kernel Ridge Regression model specification for use with the tidymodels ecosystem via **parsnip**. This spec can be paired with the "fastkrr" engine implemented in this package to fit exact or kernel approximation (Nyström, Pivoted Cholesky, Random Fourier Features) within **recipes/workflows** pipelines.

Usage

```
krr_reg(  
  mode = "regression",  
  kernel = NULL,  
  opt = NULL,  
  eps = NULL,  
  n_threads = NULL,  
  m = NULL,  
  rho = NULL,  
  penalty = NULL,  
  fastcv = NULL  
)
```

Arguments

mode	A single string; only "regression" is supported.
kernel	Kernel matrix K has two kinds of Kernel ("gaussian", "laplace").
opt	Method for constructing or approximating : "exact" Construct the full kernel matrix $K \in \mathbb{R}^{n \times n}$ using design matrix X . "nystrom" Construct a low-rank approximation of the kernel matrix $K \in \mathbb{R}^{n \times n}$ using the Nyström approximation. "pivoted" Construct a low-rank approximation of the kernel matrix $K \in \mathbb{R}^{n \times n}$ using Pivoted Cholesky decomposition. "rff" Use Random Fourier Features to construct a feature map $Z \in \mathbb{R}^{n \times m}$ (with m random features) so that $K \approx ZZ^T$. Here, m is the number of features.
eps	Tolerance parameter used only in "pivoted" for stopping criterion of the Pivoted Cholesky decomposition.
n_threads	Number of parallel threads. It is applied only for opt = "nystrom" or opt = "rff", and for the Laplace kernel (kernel = "laplace").
m	Approximation rank(number of random features) used for the low-rank kernel approximation.
rho	Scaling parameter of the kernel(ρ).
penalty	Regularization parameter.
fastcv	If TRUE, accelerated cross-validation is performed via sequential testing (early stopping) as implemented in the CVST package.

Value

A parsnip model specification of class "krr_reg".

Examples

```

if (all(vapply(
  c("parsnip", "stats", "modeldata"),
  requireNamespace, quietly = TRUE, FUN.VALUE = logical(1)
))) {
  library(tidymodels)
  library(parsnip)
  library(stats)
  library(modeldata)

  # Data analysis
  data(ames)
  ames = ames %>% mutate(Sale_Price = log10(Sale_Price))

  set.seed(502)
  ames_split = initial_split(ames, prop = 0.80, strata = Sale_Price)
  ames_train = training(ames_split) # dim (2342, 74)
  ames_test = testing(ames_split) # dim (588, 74)

```

```
# Model spec
krr_spec = krr_reg(kernel = "gaussian", opt = "exact",
                  m = 50, eps = 1e-6, n_threads = 4,
                  rho = 1, penalty = tune()) %>%
  set_engine("fastkrr") %>%
  set_mode("regression")

# Define rec
rec = recipe(Sale_Price ~ Longitude + Latitude, data = ames_train)

# workflow
wf = workflow() %>%
  add_recipe(rec) %>%
  add_model(krr_spec)

# Define hyper-parameter grid
param_grid = grid_regular(
  dials::penalty(range = c(-10, -3)),
  levels = 5
)

# CV setting
set.seed(123)
cv_folds = vfold_cv(ames_train, v = 5, strata = Sale_Price)

# Tuning
tune_results = tune_grid(
  wf,
  resamples = cv_folds,
  grid = param_grid,
  metrics = metric_set(rmse),
  control = control_grid(verbose = TRUE, save_pred = TRUE)
)

# Result check
collect_metrics(tune_results)

# Select best parameter
best_params = select_best(tune_results, metric = "rmse")

# Finalized model spec using best parameter
final_spec = finalize_model(krr_spec, best_params)
final_wf = workflow() %>%
  add_recipe(rec) %>%
  add_model(final_spec)

# Finalized fitting using best parameter
final_fit = final_wf %>% fit(data = ames_train)

# Prediction
predict(final_fit, new_data = ames_test)
print(best_params)
```

```
}

```

```
make_kernel
```

```
Kernel matrix K construction for given datasets
```

Description

Constructs a kernel matrix $K \in \mathbb{R}^{n \times n'}$ given two datasets $X \in \mathbb{R}^{n \times d}$ and $X' \in \mathbb{R}^{n' \times d}$, where $x_i \in \mathbb{R}^d$ and $x'_j \in \mathbb{R}^d$ denote the i-th and j-th rows of X and X' , respectively, and $K_{ij} = \mathcal{K}(x_i, x'_j)$ for a user-specified kernel. Implemented in C++ via RcppArmadillo.

Arguments

X	Design matrix $X \in \mathbb{R}^{n \times d}$ (rows $x_i \in \mathbb{R}^d$).
X_new	Second matrix $X' \in \mathbb{R}^{n' \times d}$ (rows $x'_j \in \mathbb{R}^d$). If omitted, $X' = X$ and $n' = n$.
kernel	Kernel type; one of "gaussian" or "laplace".
rho	Kernel width parameter ($\rho > 0$).
n_threads	Number of parallel threads. The default is 4. If the system does not support 4 threads, it automatically falls back to 1 thread. Parallelization (implemented in C++) is one of the main advantages of this package and is applied only for "laplace" kernels.

Details

Gaussian:

$$\mathcal{K}(x_i, x_j) = \exp(-\rho \|x_i - x_j\|_2^2)$$

Laplace:

$$\mathcal{K}(x_i, x_j) = \exp(-\rho \|x_i - x_j\|_1)$$

Value

An S3 object of class "kernel_matrix" that represents the computed kernel matrix. If X_new is NULL, the result is a symmetric matrix $K_{ij} = \mathcal{K}(x_i, x_j)$, with $K \in \mathbb{R}^{n \times n}$. Otherwise, the result is a rectangular matrix $K'_{ij} = \mathcal{K}(x_i, x'_j)$, with $K' \in \mathbb{R}^{n \times n'}$.

Examples

```
# Data setting
set.seed(1)
d = 1
rho = 1
n = 1000
X = matrix(runif(n*d, 0, 1), nrow = n, ncol = d)

# New design matrix
new_n = 1500
```

```

new_X = matrix(runif(new_n*d, 0, 1), nrow = new_n, ncol = d)

# Make kernel : Gaussian kernel
K = make_kernel(X, kernel = "gaussian", rho = rho) ## symmetric matrix
new_K = make_kernel(X, new_X, kernel = "gaussian", rho = rho) ## rectangular matrix

# Make kernel : Laplace kernel
K = make_kernel(X, kernel = "laplace", rho = rho, n_threads = 1) ## symmetric matrix
new_K = make_kernel(X, new_X, kernel = "laplace", rho = rho, n_threads = 1) ## rectangular matrix

```

param

Extract/print hyperparameters of fitted models

Description

"param()" is a generic S3 function that displays (and invisibly returns) model hyperparameters. Methods are provided for "krr" objects.

Usage

```

param(x, ...)

## Default S3 method:
param(x, ...)

```

Arguments

x An object.
... Additional arguments passed to methods.

Value

A named list of hyperparameters (invisibly); may print side effects.

param.krr

Param method for fitted Kernel Ridge Regression models

Description

Displays (and invisibly returns) the hyperparameters actually used by a fitted object. For krr objects returned by `fastkrr`, this prints a concise hyperparameter panel (e.g., rho, lambda, m, eps, n_threads, d).

Usage

```
## S3 method for class 'krr'  
param(x, ...)
```

Arguments

`x` An object of class "krr", typically returned by [fastkrr](#).

`...` Additional arguments.

Details

Pivoted approximation note: When `opt = "pivoted"`, the effective number of pivots `m` used during the approximation may be smaller than the user-specified `m` because the algorithm can stop early based on `eps`. If you want to confirm the initial `m` that you set, please see the printed *Call* (the original function call shows your input arguments).

Value

Prints a human-readable panel and returns (invisibly) a named list of hyperparameters.

See Also

[fastkrr](#)

Examples

```
# Data setting  
set.seed(1)  
lambda = 1e-4  
d = 1  
n = 50  
rho = 1  
X = matrix(runif(n*d, 0, 1), nrow = n, ncol = d)  
y = as.vector(sin(2*pi*rowMeans(X)^3) + rnorm(n, 0, 0.1))  
  
model = fastkrr(X, y, kernel="gaussian", opt="nystrom",  
               rho=1, lambda=1e-4, m=200, n_threads=4, fastcv=FALSE)  
  
class(model)  
param(model)
```

`plot.krr`*Plot method for fitted Kernel Ridge Regression (KRR) models*

Description

Visualizes fitted results from a Kernel Ridge Regression (KRR) model. Automatically generates predictions on a regular grid (120% of training sample size) and overlays them with training data.

Usage

```
## S3 method for class 'krr'  
plot(x, show_points = TRUE, ...)
```

Arguments

<code>x</code>	A fitted KRR model (class "krr") returned by fastkrr .
<code>show_points</code>	Logical; if TRUE, displays the training data points. Default = TRUE.
<code>...</code>	Additional arguments (currently ignored).

Details

For multivariate inputs ($d \geq 2$), visualization requires fixing all but one variable. For example, in 2D, one can plot $f(x_1, x_2 = \bar{x}_2)$ to examine the effect of x_1 while holding x_2 at its mean.

Value

A ggplot object showing the fitted regression curve.

See Also

[fastkrr](#), [predict.krr](#)

Examples

```
set.seed(1)  
n = 1000  
rho = 1  
X = runif(n, 0, 1)  
y = sin(2*pi*X^3) + rnorm(n, 0, 0.1)  
  
model_exact = fastkrr(X, y, kernel = "gaussian", rho = rho, opt = "exact", verbose = FALSE)  
plot(model_exact)
```

predict.krr	<i>Predict responses for new data using fitted KRR model</i>
-------------	--

Description

Generates predictions from a fitted Kernel Ridge Regression (KRR) model for new data.

Usage

```
## S3 method for class 'krr'
predict(object, newdata, ...)
```

Arguments

object	A S3 object of class krr created by fastkrr .
newdata	New design matrix or data frame containing new observations for which predictions are to be made. If newdata is missing, the function returns fitted values.
...	Additional arguments (currently ignored).

Value

A numeric vector of predicted values corresponding to newdata or fitted values.

See Also

[fastkrr](#), [make_kernel](#)

Examples

```
# Data setting
n = 30
d = 1
X = matrix(runif(n*d, 0, 1), nrow = n, ncol = d)
y = as.vector(sin(2*pi*rowMeans(X)^3) + rnorm(n, 0, 0.1))
lambda = 1e-4
rho = 1

# Fitting model: pivoted
model = fastkrr(X, y, kernel = "gaussian", rho = rho, lambda = lambda, opt = "pivoted")

# Predict
new_n = 50
new_x = matrix(runif(new_n*d, 0, 1), nrow = new_n, ncol = d)
new_y = as.vector(sin(2*pi*rowMeans(new_x)^3) + rnorm(new_n, 0, 0.1))

pred = predict(model, new_x)
crossprod(pred - new_y) / new_n

predict(model) == attributes(model)$fitted.values
```

```
print.approx_kernel
```

Print method for approximated kernel matrices

Description

Displays the approximated kernel matrix and key options used to construct it.

Usage

```
## S3 method for class 'approx_kernel'  
print(x, ...)
```

Arguments

`x` An S3 object created by [approx_kernel](#).
`...` Additional arguments (currently ignored).

Details

The function prints the stored approximated kernel matrix (top-left 6x6) and summarizes options such as the approximation method (`opt`), approximation degree (`m`), numerical tolerance (`eps`), and number of threads used (`n_threads`).

Value

An approximated kernel matrix and its associated options.

See Also

[approx_kernel](#), [print.krr](#), [print.kernel_matrix](#)

Examples

```
# Data setting  
set.seed(1)  
d = 1  
n = 1000  
m = 50  
rho = 1  
X = matrix(runif(n*d, 0, 1), nrow = n, ncol = d)  
y = as.vector(sin(2*pi*rowMeans(X)^3) + rnorm(n, 0, 0.1))  
K = make_kernel(X, kernel = "gaussian", rho = rho)  
  
# Example: nystrom  
K_nystrom = approx_kernel(K = K, opt = "nystrom", m = m, d = d, rho = rho, n_threads = 1)  
class(K_nystrom)  
  
print(K_nystrom)
```

print.kernel_matrix *Print method for kernel matrices*

Description

Displays the top-left 6x6 portion of a kernel or approximated kernel matrix for quick inspection.

Usage

```
## S3 method for class 'kernel_matrix'  
print(x, ...)
```

Arguments

x An object of class `kernel_matrix`, which may represent either an exact kernel matrix (from [make_kernel](#) or [fastkrr](#)) or an approximated kernel matrix (from [approx_kernel](#)).

... Additional arguments (currently ignored).

Value

A top-left 6x6 block of the kernel matrix to the console.

See Also

[approx_kernel](#), [fastkrr](#), [print.approx_kernel](#), [print.krr](#)

Examples

```
# data setting  
set.seed(1)  
n = 1000 ; d = 1  
m = 100  
rho = 1  
X = matrix(runif(n*d, 0, 1), nrow = n, ncol = d)  
y = as.vector(sin(2*pi*X^3) + rnorm(n, 0, 0.1))  
  
# Example for fastkrr  
fit_pivoted = fastkrr(X, y,  
                      kernel = "gaussian", opt = "pivoted",  
                      m = 100, fastcv = TRUE, verbose = FALSE)  
  
class(attr(fit_pivoted, "K"))  
print(class(attr(fit_pivoted, "K")))  
  
class(attr(fit_pivoted, "K_approx"))  
print(class(attr(fit_pivoted, "K_approx")))
```

```

# Example for make_kernel
K = make_kernel(X, kernel = "gaussian", rho = rho)

class(K)
print(K)

# Example for make_kernel
K_rff = approx_kernel(X = X, opt = "rff", kernel = "gaussian",
                     d = d, rho = rho, n_threads = 1, m = 100)

class(attr(K_rff, "K_approx"))
print(attr(K_rff, "K_approx"))

```

print.krr

Print method for fitted Kernel Ridge Regression models

Description

Displays key information from a fitted Kernel Ridge Regression (KRR) model, including the original call, first few coefficients, a 6×6 block of the kernel (or approximated kernel) matrix, and the main kernel options.

Usage

```
## S3 method for class 'krr'
print(x, ...)
```

Arguments

`x` An S3 object of class `krr`, typically returned by [fastkrr](#).
`...` Additional arguments (currently ignored).

Value

A human-readable summary of the fitted KRR model to the console.

See Also

[fastkrr](#), [print.approx_kernel](#), [print.kernel_matrix](#)

Examples

```

# Data setting
set.seed(1)
lambda = 1e-4
d = 1
n = 50

```

```

rho = 1
X = matrix(runif(n*d, 0, 1), nrow = n, ncol = d)
y = as.vector(sin(2*pi*rowMeans(X)^3) + rnorm(n, 0, 0.1))

# Example: exact
model = fastkrr(X, y,
               kernel = "gaussian", opt = "exact",
               rho = rho, lambda = 1e-4)

class(model)

print(model)

```

summary.krr

*Summary method for fitted Kernel Ridge Regression models***Description**

Displays key information from a fitted Kernel Ridge Regression (KRR) model, including the original call, first few coefficients, a 6×6 block of the kernel (or approximated kernel) matrix, and the main kernel options.

Usage

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

Arguments

`object` An S3 object of class `krr`, typically returned by `fastkrr`.
`...` Additional arguments (currently ignored).

Value

A human-readable summary of the fitted KRR model to the console.

See Also

[fastkrr](#)

Examples

```

# Data setting
set.seed(1)
lambda = 1e-4
d = 1
n = 50
rho = 1
X = matrix(runif(n*d, 0, 1), nrow = n, ncol = d)
y = as.vector(sin(2*pi*rowMeans(X)^3) + rnorm(n, 0, 0.1))

```

```
# Example: exact
model = fastkrr(X, y,
               kernel = "gaussian", opt = "exact",
               rho = rho, lambda = 1e-4)

class(model)

summary(model)
```

tunable.krr_reg *Expose tunable parameters for "krr_reg"*

Description

Supplies a tibble of tunable arguments for "krr_reg".

Usage

```
## S3 method for class 'krr_reg'
tunable(x, ...)
```

Arguments

x A "krr_reg" model specification.
... Not used; included for S3 method compatibility.

Value

A tibble (one row per tunable parameter) with columns "name", "call_info", "source", "component", and "component_id".

Index

* package

FastKRR-package, 2

approx_kernel, 3, 20, 21

coef.krr, 6

error, 7

error.krr, 8

FastKRR (FastKRR-package), 2

fastkrr, 6, 8, 9, 16–19, 21–23

FastKRR-package, 2

krr_reg, 12

make_kernel, 15, 19, 21

param, 16

param.krr, 16

plot.krr, 8, 18

predict.krr, 8, 18, 19

print.approx_kernel, 20, 21, 22

print.kernel_matrix, 20, 21, 22

print.krr, 20, 21, 22

summary.krr, 8, 23

tunable.krr_reg, 24