

Package ‘RcppEnsmallen’

May 7, 2026

Title Header-Only C++ Mathematical Optimization Library for
'Armadillo'

Version 0.3.10.0.1

Description 'Ensmallen' is a templated C++ mathematical optimization library (by the 'MLPACK' team) that provides a simple set of abstractions for writing an objective function to optimize. Provided within are various standard and cutting-edge optimizers that include full-batch gradient descent techniques, small-batch techniques, gradient-free optimizers, and constrained optimization. The 'RcppEnsmallen' package includes the header files from the 'Ensmallen' library and pairs the appropriate header files from 'armadillo' through the 'RcppArmadillo' package. Therefore, users do not need to install 'Ensmallen' nor 'Armadillo' to use 'RcppEnsmallen'. Note that 'Ensmallen' is licensed under 3-Clause BSD, 'Armadillo' starting from 7.800.0 is licensed under Apache License 2, 'RcppArmadillo' (the 'Rcpp' bindings/bridge to 'Armadillo') is licensed under the GNU GPL version 2 or later. Thus, 'RcppEnsmallen' is also licensed under similar terms. Note that 'Ensmallen' requires a compiler that supports 'C++14' and 'Armadillo' 10.8.2 or later.

Depends R (>= 4.3.0)

License GPL (>= 2)

URL <https://github.com/coatless-rpkg/rcppensmallen>,
<https://r-pkg.thecoatlessprofessor.com/rcppensmallen/>,
<https://github.com/mlpack/ensmallen>, <https://ensmallen.org/>

BugReports <https://github.com/coatless-rpkg/rcppensmallen/issues>

Encoding UTF-8

LinkingTo Rcpp, RcppArmadillo

Imports Rcpp

RoxygenNote 7.3.3

Suggests knitr, rmarkdown

VignetteBuilder knitr

NeedsCompilation yes

Author James Joseph Balamuta [aut, cre, cph] (ORCID:
 <<https://orcid.org/0000-0003-2826-8458>>),
 Dirk Eddelbuettel [aut, cph] (ORCID:
 <<https://orcid.org/0000-0001-6419-907X>>)

Maintainer James Joseph Balamuta <balamut2@illinois.edu>

Repository CRAN

Date/Publication 2025-09-30 22:10:02 UTC

Contents

RcppEnsmallen-package	2
lin_reg_lbfgs	3
Index	5

RcppEnsmallen-package *RcppEnsmallen: Header-Only C++ Mathematical Optimization Library for 'Armadillo'*

Description

'Ensmallen' is a templated C++ mathematical optimization library (by the 'MLPACK' team) that provides a simple set of abstractions for writing an objective function to optimize. Provided within are various standard and cutting-edge optimizers that include full-batch gradient descent techniques, small-batch techniques, gradient-free optimizers, and constrained optimization. The 'RcppEnsmallen' package includes the header files from the 'Ensmallen' library and pairs the appropriate header files from 'armadillo' through the 'RcppArmadillo' package. Therefore, users do not need to install 'Ensmallen' nor 'Armadillo' to use 'RcppEnsmallen'. Note that 'Ensmallen' is licensed under 3-Clause BSD, 'Armadillo' starting from 7.800.0 is licensed under Apache License 2, 'RcppArmadillo' (the 'Rcpp' bindings/bridge to 'Armadillo') is licensed under the GNU GPL version 2 or later. Thus, 'RcppEnsmallen' is also licensed under similar terms. Note that 'Ensmallen' requires a compiler that supports 'C++14' and 'Armadillo' 10.8.2 or later.

Author(s)

Maintainer: James Joseph Balamuta <balamut2@illinois.edu> (ORCID) [copyright holder]

Authors:

- Dirk Eddelbuettel <edd@debian.org> (ORCID) [copyright holder]

See Also

Useful links:

- <https://github.com/coatless-rpkg/rcppensmallen>
- <https://r-pkg.thecoatlessprofessor.com/rcppensmallen/>

- <https://github.com/mlpack/enmallen>
- <https://enmallen.org/>
- Report bugs at <https://github.com/coatless-rpkg/rcppenmallen/issues>

`lin_reg_lbfgs`*Linear Regression with L-BFGS*

Description

Solves the Linear Regression's Residual Sum of Squares using the L-BFGS optimizer.

Usage

```
lin_reg_lbfgs(X, y)
```

Arguments

`X` A matrix that is the Design Matrix for the regression problem.
`y` A vec containing the response values.

Details

Consider the **Residual Sum of Squares**, also known as **RSS**, defined as:

$$RSS(\beta) = (\mathbf{y} - \mathbf{X}\beta)^T (\mathbf{y} - \mathbf{X}\beta)$$

The objective function is defined as:

$$f(\beta) = (y - X\beta)^2$$

The gradient is defined as:

$$\frac{\partial RSS}{\partial \beta} = -2\mathbf{X}^T (\mathbf{y} - \mathbf{X}\beta)$$

Value

The estimated β parameter values for the linear regression.

Examples

```
# Number of Points
n = 1000

# Select beta parameters
beta = c(-2, 1.5, 3, 8.2, 6.6)

# Number of Predictors (including intercept)
p = length(beta)
```

```
# Generate predictors from a normal distribution
X_i = matrix(rnorm(n), ncol = p - 1)

# Add an intercept
X = cbind(1, X_i)

# Generate y values
y = X%%beta + rnorm(n / (p - 1))

# Run optimization with lbfgs
theta_hat = lin_reg_lbfgs(X, y)

# Verify parameters were recovered
cbind(actual = beta, estimated = theta_hat)
```

Index

`lin_reg_lbfgs`, [3](#)

`RcppEnsmallen (RcppEnsmallen-package)`, [2](#)

`RcppEnsmallen-package`, [2](#)