

Package: l1ball (via r-universe)

August 25, 2024

Type Package

Title L1-Ball Prior for Sparse Regression

Version 0.1.0

Author Maoran Xu and Leo L. Duan

Maintainer Maoran Xu <maoranxu@ufl.edu>

Description Provides function for the l1-ball prior on high-dimensional regression. The main function, `l1ball()`, yields posterior samples for linear regression, as introduced by Xu and Duan (2020) <[arXiv:2006.01340](https://arxiv.org/abs/2006.01340)>.

License GPL (>= 2)

Depends R (>= 3.1.0)

Imports VGAM, stats

Encoding UTF-8

LazyData true

RoxygenNote 7.1.0

NeedsCompilation no

Repository CRAN

Date/Publication 2020-07-24 17:10:02 UTC

Contents

l1ball	2
Index	3

11ball

*Fit the L1 prior***Description**

This package provides an implementation of the Gibbs sampler, for using 11-ball prior with the regression likelihood $y_i = X_i\theta + \epsilon_i$, $\epsilon_i \sim N(0, \sigma^2)$.

Arguments

y	A data vector, n by 1
X	A design matrix, n by p
b_w	The parameter in $Beta(1, p^{b_w})$ for w , default $b_w = 1$
step	Number of steps to run the Markov Chain Monte Carlo
burnin	Number of burn-ins
b_lam	The parameter in $\lambda_i \sim Inverse - Gamma(1, b_\lambda)$, default $b_\lambda = 10^{-3}$. To increase the level of shrinkage, use smaller b_λ .

Value

The posterior sample collected from the Markov Chain:

- trace_theta: θ
- trace_NonZero: The non-zero indicator $1(\theta_i \neq 0)$
- trace_Lam: λ_i
- trace_Sigma: σ^2

Examples

```
n = 60
p = 100
X <- matrix(rnorm(n*p), n, p)
d = 5
w0 <- c(rep(0, p-d), rnorm(d)*0.1+1)
y = X%% w0 + rnorm(n, 0, .1)
trace <- 11ball(y, X, steps=2000, burnin = 2000)
plot(colMeans(trace$trace_theta))
```

Index

11ball, 2